

MORSE HOMOLOGY, TROPICAL GEOMETRY, AND HOMOLOGICAL MIRROR SYMMETRY FOR TORIC VARIETIES

MOHAMMED ABOUZAIID

ABSTRACT. Given a smooth projective toric variety X , we construct an A_∞ category of Lagrangians with boundary on a level set of the Landau-Ginzburg mirror of X . We prove that this category is quasi-equivalent to the DG category of line bundles on X .

CONTENTS

1. Introduction	1
2. Outline and discussion of the proof	6
3. The category of tropical Lagrangian sections	9
4. General framework	17
5. HMS for toric varieties: Fukaya = Morse	35
6. HMS for toric varieties: Morse = Čech	40
Appendix A. A second look at semi-tropical degenerations	46
Appendix B. Moduli spaces of trees	57
Appendix C. Orientations in Morse theory	61
Appendix D. A_∞ pre-categories	67
Appendix E. Cellular chains and simplicial cochains on manifolds with boundary	69
References	73

1. INTRODUCTION

Mirror symmetry entered the mathematical realm when, in [7], physicists were able to enumerate curves on the quintic threefold by studying the complex moduli space of an algebraic variety which is most often referred to as the “mirror quintic.” The fact that such a computation was impossible using then available mathematical techniques contributed to the development of an entirely new field of mathematics (Gromov-Witten theory), which eventually resulted in Givental’s proof in [19, 20] of a mirror conjecture for complete intersections in toric varieties. One important feature apparent in Givental’s work is that, whereas the mirror of a Calabi-Yau variety (one whose canonical bundle is trivial) is another Calabi-Yau variety, the mirrors of Fano varieties (essentially, those with positive anti-canonical bundle) are Landau-Ginzburg models; i.e. affine varieties equipped with a map to \mathbb{C} called the superpotential.

Meanwhile, there have been attempts at explaining the mirror phenomenon. Most relevant to us are categorical ideas of Kontsevich and geometric ideas of

Strominger, Yau, Zaslow. In his 1994 ICM address, Kontsevich sought to explain mirror symmetry for pairs of mirror projective Calabi-Yau manifolds M and \tilde{M} as a consequence of an equivalence of categories which he called homological mirror symmetry:

HMS Conjecture for CY ([26]). *The derived category of coherent sheaves on M is equivalent to the derived Fukaya category of \tilde{M} .*

A few words of explanations are in order. Assuming that M is smooth, the derived category $D^b \text{Coh}(M)$ of coherent sheaves on M is generated by (algebraic or, equivalently, holomorphic) vector bundles on M and the morphisms in this category are essentially determined by the higher cohomology groups of such vector bundles. In particular, $D^b \text{Coh}(M)$ depends only on the complex structure of M , and not on any projective embedding.

On the other hand, the derived Fukaya category depends on the symplectic form (essentially the imaginary part of the Kähler form) induced by a projective embedding, and is independent of the complex structure. In essence, the study of the Fukaya category is a problem in symplectic topology, and not in algebraic geometry. Although current technology makes a satisfactory definition of the Fukaya category possible, setting up all the preliminaries would take us far beyond what would be appropriate for an introduction. Let us simply say that the objects of such a category are Lagrangian submanifolds (i.e. manifolds on which the symplectic form vanishes) and that the definition of the spaces of morphisms and of the composition entails the study of holomorphic curves with respect to an almost complex structure which is compatible with the symplectic form. Lest the reader object that the Fukaya category is supposed to depend only the symplectic structure, one should note that, while the choice of an almost complex structure is necessary to give a definition, the quasi-isomorphism class of the resulting category is independent of this choice. This is essentially the same phenomenon which makes the Gromov-Witten invariants of an algebraic variety deformation invariants.

In his 1999 ENS lectures [27], Kontsevich extended his homological mirror symmetry to Fano varieties. Since this is the main subject of this paper, we shall return to it in more detail. Before continuing with homological mirror symmetry, let us discuss the SYZ explanation of mirror symmetry. As is clear from the title of their paper [42], their main conjecture is:

SYZ Conjecture. *If M and \tilde{M} are a mirror pair of Calabi-Yau threefolds, then there exists (singular) special Lagrangian torus fibrations $M \rightarrow B$ and $\tilde{M} \rightarrow B$ which are dual to each other.*

To understand what is meant by dual fibrations, it is best to consider an example which generalizes to describe the phenomenon whenever there are no singularities. Consider any smooth manifold B . It is well known that the cotangent bundle T^*B is a symplectic manifold, whereas an infinitesimal neighbourhood of the zero section in the tangent bundle TB admits a canonical complex structure. In nice situations, this complex structure extends to the entire tangent space, and for simplicity let us assume that this is the case. Upon choosing a metric on B , we induce a complex structure on T^*B and a symplectic structure on TB . In order to obtain torus fibrations, one should consider a manifold equipped with an integral affine structure, i.e. a reduction of the structure group of B from $\text{Diff}(\mathbb{R}^n)$ to the group of diffeomorphisms of \mathbb{R}^n whose differentials preserve the standard lattice.

Such a manifold comes equipped with lattices in its tangent and cotangent bundles. The corresponding quotients are the desired (prototypical) dual torus fibrations. For discussion of mirror symmetry for non-singular torus fibrations, see [29].

The only part of the SYZ conjecture relevant to this paper is the following important observation which is by now folkloric:

Conjecture 1. *The mirror of a Lagrangian section of $M \rightarrow B$ is a holomorphic line bundle on \check{M} .*

While the SYZ conjecture has only been stated for Calabi-Yau manifolds, we shall nonetheless use the intuition that it provides. We note that toric varieties are, in a sense, equipped with Lagrangian torus fibrations, so the appearance of ideas from SYZ in a discussion of homological mirror symmetry for toric varieties should not be surprising.

Statement of results. The main result concerns half of Kontsevich's homological mirror symmetry conjecture:

HMS Conjecture for Fano ([27]). *If the anti-canonical bundle of X is ample (i.e. X is Fano), the bounded derived category of coherent sheaves on X , $D^b \text{Coh}(X)$, is equivalent, as an A_∞ category, to the Karoubi completion of the derived Fukaya category of the mirror of X .*

The conjecture does not include a description of this conjectural mirror, although a wide class of mirrors is described in [24]. In fact, we will be working with smooth projective toric varieties even when their anti-canonical bundles are not ample. However, as we shall see in the discussion following the statement of Theorem 1.2, we cannot expect an equivalence of categories without the Fano condition.

The fact that the homological mirror symmetry conjecture is a statement about A_∞ categories means that we are seeking equivalences at the chain level. The topologically inclined reader should have in mind secondary products on cohomology (for example the Massey products) which could not be detected by someone who simply thinks of cohomology as a functor satisfying the Eilenberg-Steenrod axioms; these are essentially chain-level phenomena. The algebraically inclined reader should have in mind the fact that derived categories, while extremely useful, have been increasingly supplanted in recent years by DG categories, which remember more information about the original object which is often lost by deriving, see [8]. In an appropriate homotopy category, every (unital) A_∞ category is equivalent to a DG category.

Let us briefly describe what we mean by the derived Fukaya category of the mirror of X . Recall from the first part of the introduction that the mirrors of Fano varieties are superpotentials

$$W: Y \rightarrow \mathbb{C}.$$

Kontsevich suggested that correct notion of a Fukaya category which takes into account the superpotential has, as objects, Lagrangian submanifolds $L \subset Y$ whose image under W agrees with a line $\gamma: [0, \infty) \rightarrow \mathbb{C}$ away from a compact subset. There is no loss of generality in assuming that the curve γ starts at the origin. Assuming that W is a symplectic Lefschetz fibration (this is a generic situation), it is not hard to see that the subset of L which lies over γ can be obtained by parallel transport, along γ , of the intersection of L with $W^{-1}(0)$. In particular, we lose no

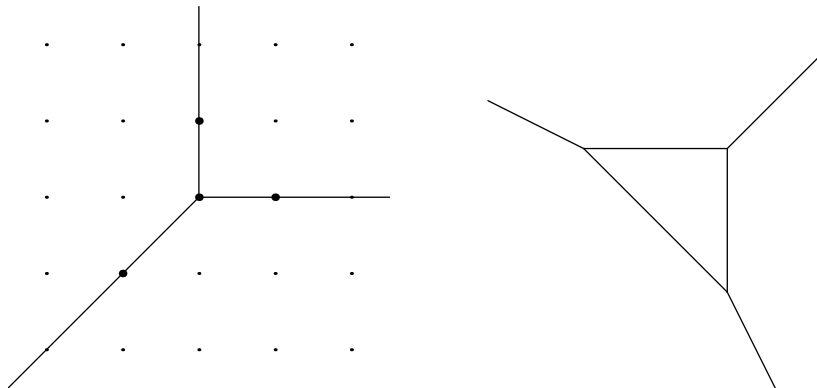


FIGURE 1. The fan defining \mathbb{CP}^2 and a corresponding tropical amoeba.

information if we work with compact Lagrangians whose boundary lies on $W^{-1}(0)$. This is the point of view taken in this paper.

Remark 1.1. Defining the Fukaya category using only holomorphic curves (and no perturbations) does not yield a category because of issues of transversality. If one is willing to make choices of inhomogeneous perturbations, it is often possible to extract an honest A_∞ category from the geometry of the Fukaya category. We make an alternative choice in this paper, working with the notion of pre-categories, which we briefly discuss in Appendix D. The same remark applies to the Morse pre-categories which we shall be using.

From now on, we consider the case where X is a smooth projective toric variety. The mirror of such a variety is expected to be a Laurent polynomial W on $(\mathbb{C}^*)^n$ which can be explicitly obtained from the fan Δ which defines X as follows:

Consider the 1-cones of Δ . These are half rays starting at the origin which have rational slope, so they intersect the lattice at a unique primitive (non-divisible) vector. Let A be the set which consists of the primitive vertices of these 1-cones and of the origin. The mirror of X is the family of Laurent polynomials

$$\sum_{\alpha \in A} c_\alpha z^\alpha,$$

where $z^\alpha = z^{a_1} \dots z^{a_n}$ whenever $\alpha = (a_1, \dots, a_n)$. From our point of view the specific coefficients c_α are irrelevant as long as they are chosen generically; the corresponding pairs $((\mathbb{C}^*)^n, W^{-1}(0))$ are symplectomorphic. In [1] we consider a family of Lefschetz fibrations $W_{t,s}$ which are not given by Laurent polynomials for $s \neq 0$, but for which the pairs $((\mathbb{C}^*)^n, W_{t,s}^{-1}(0))$ are still symplectomorphic for different values of s and t . The purpose of these deformations away from Laurent polynomials is to obtain precise control for the convergence of the hypersurfaces to certain piecewise linear symplectic hypersurfaces constructed from tropical geometry. From now on, we write M for any of these hypersurfaces $W_{t,s}^{-1}(0)$.

Given such a pair, we define a (relative) Fukaya pre-category $\text{Fuk}((\mathbb{C}^*)^n, M)$ consisting of compact Lagrangians on $(\mathbb{C}^*)^n$ whose boundary lies in M as described above. Some details are given in Section 3.1. In this paper, we prove

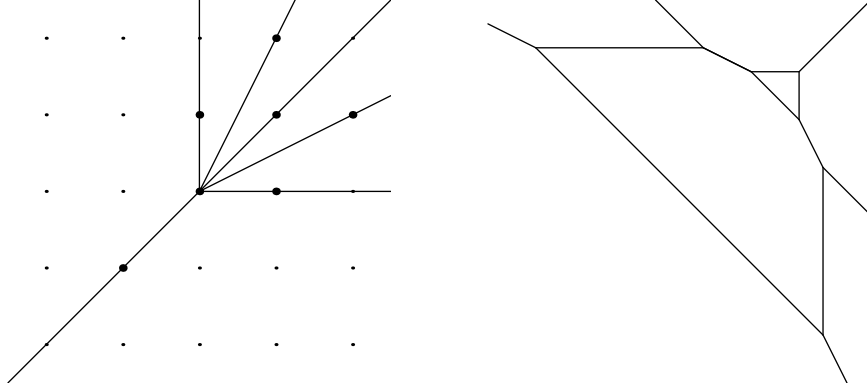


FIGURE 2. The fan defining an iterated blowup of \mathbb{CP}^2 and a corresponding tropical amoeba.

Theorem 1.2. *There exists a full sub-pre-category of $\mathrm{Fuk}((\mathbb{C}^*)^n, M)$ which is quasi-equivalent as an A_∞ pre-category to the category of line bundles on X .*

In order for the statement to make sense, we can equip the category of line bundle with the structure of a *DG* category by using the (canonical) cover by affine toric subvarieties and passing through Čech cohomology in order to compute the cohomology of holomorphic line bundles.

As one might expect from Conjecture 1 the relevant sub-pre-category of $\mathrm{Fuk}((\mathbb{C}^*)^n, M)$ is obtained by studying Lagrangian sections of a torus fibration

$$(\mathbb{C}^*)^n \rightarrow \mathbb{R}^n.$$

However, the sections are only taken over a subset of the base \mathbb{R}^n . For example, the Lagrangians which are mirror to line bundles on \mathbb{CP}^2 are essentially sections over the triangle which appears in the picture on the right of Figure 1. The reader might recognize this triangle as the moment polytope of \mathbb{CP}^2 , but this is not how it appears in our construction. Rather, following Mikhalkin [33], the graph on the right of Figure 1 is the tropical amoeba of the one parameter family of Laurent polynomials $W_t(x, y) = 1 + t^{-1}x + t^{-1}y + \frac{t^{-1}}{xy}$ which is the Landau-Ginzburg mirror of \mathbb{CP}^2 .

The differences with Kontsevich's conjecture are that we prove the result before passing to any derived category and that the toric variety X is not required to have ample anti-canonical bundle. As Figure 2 illustrates, one of the consequences of dropping the Fano condition is that, while the moment polytope of the smooth toric variety still appears as a component of the complement of the tropical amoeba, there may be other bounded polytopes. In Figure 2, Lagrangian sections over the small triangle form a sub-pre-category of $\mathrm{Fuk}((\mathbb{C}^*)^n, M)$ whose objects do not correspond to any line bundles on X (in fact, with a bit more work, one can see that such sections do not correspond to any object of $D^b \mathrm{Coh}(X)$).

In any case, Seidel observed that one can use a theorem of Thomason to prove the following result:

Proposition 1.3. *If X is a smooth toric variety, then $D^b \mathrm{Coh}(X)$ is generated by line bundles.*

Proof. In [43, Theorem 2.1], Thomason proves that there is a bijective correspondence between subgroups of $K_0(X)$ (the algebraic K -theory of X), and full triangulated subcategories \mathcal{A} of $D^b \text{Coh}(X)$ whose idempotent closure is $D^b \text{Coh}(X)$ (i.e. such that every object of $D^b \text{Coh}(X)$ is a summand of an object in \mathcal{A}). By an observation of Kontsevich, the full triangulated category generated by line bundles satisfies this second condition if X is smooth. Indeed, we know that every coherent sheaf has an (infinite) resolution by direct sums of line bundles. Consider the truncation of such a resolution at the $n + 1$ 'st position. Since the category of coherent sheaves on a smooth variety of dimension n has cohomological dimension n , there are no non-trivial extensions of length greater than n , so that this extension splits (in the derived category).

To complete the proof of the proposition, it suffices to know that $K_0(X)$ is generated by line bundles. This is a classical result in toric geometry. A stronger result on equivariant K -theory of toric varieties is, for example, proved in [35, Proposition 3]. \square

We can therefore use Theorem 1.2 to conclude

Corollary 1.4. *$D^b \text{Coh}(X)$ is equivalent to a full subcategory of $D^\pi \text{Fuk}((\mathbb{C}^*)^n, M)$.* \square

Remark 1.5. This paper is essentially a continuation of [1] whose results and techniques we use extensively. The major difference is that, in [1], we were able to explicitly identify certain moduli spaces of discs when we restricted ourselves to a fixed ample line bundle and its powers, whereas we are unable to do this in the general case. Instead, as explained in the next section, we rely on the equivalence between Floer and Morse homology in order to prove the results of this paper.

The expert reader will notice that we do not rely on the notion of vanishing cycles of Lefschetz fibrations to construct objects of the Fukaya category. This distinguishes our approach from that taken by Auroux, Katzarkov, and Orlov [2, 3] in their proof of the mirror conjecture for Fano surfaces. On the other hand, using techniques somewhat closer to ours in the sense that they rely on Morse theory, Bondal and Ruan have announced a similar result for weighted projective spaces, see [6]. More recently, Fang, Liu, Treumann, and Zaslow [11] have given a different proof of Homological Mirror symmetry relying on the correspondence established by Nadler and Zaslow between constructible sheaves on the base of a compact manifold and Lagrangians in its cotangent bundle. Their approach is in a sense dual to the one we've taken.

Acknowledgments. I would first like to thank Paul Seidel for guidance and encouragement. Conversations with Mitya Boyarchenko, Ludmil Katzarkov, Conan Leung, Peter May, and Ivan Smith helped me connect ideas together and avoid dangerous pitfalls. Discussions with Kevin Costello, and his comments on an early draft, were particularly helpful. Finally, I received extensive and helpful comments from Denis Auroux, Mark Gross, and an anonymous referee.

This work was supported by the 2006 Highbridge Research Prize and was partially conducted during the period the author served as Clay Research Fellow.

2. OUTLINE AND DISCUSSION OF THE PROOF

In logical order, the proof proceeds as follows:

- (1) Identify the appropriate sub-pre-category of $\mathrm{Fuk}((\mathbb{C}^*)^n, M)$.
- (2) Prove that this sub-pre-category is quasi-isomorphic to a pre-category with essentially the same objects, but whose morphisms can be computed by studying Morse theory on the moment polytope Q of X . We denote this pre-category $\mathrm{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$, where \mathcal{Q} is described below.
- (3) Prove that the pre-category $\mathrm{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$ is itself equivalent to a pre-category $\mathrm{Simp}^{\mathbb{Z}^n}(Q_b)$ whose morphisms are certain subcomplexes of the cochain complex of Q_b , the barycentric subdivision of Q .
- (4) Identify $\mathrm{Simp}^{\mathbb{Z}^n}(Q_b)$ with the category $\check{\mathrm{Cech}}(X)$ of line bundles on X .

Implicit in the above description is the fact that each of the maps we construct are in fact A_∞ functors. Let us forget for a moment about the A_∞ structure, and remember only the structures at the level of homology. The previous steps become

- (1) The appropriate sub-pre-category of $\mathrm{Fuk}((\mathbb{C}^*)^n, M)$ was already implicit in [1], particularly in Section 5. It essentially consists of sections of a torus bundle over a smoothing \mathcal{Q} of the moment polytope, with appropriate boundary conditions. The universal cover of this symplectic manifold is the cotangent bundle of \mathcal{Q} .
- (2) Floer proved in [12] that Morse homology agrees with Floer homology of exact sections of the cotangent bundle. Extending the result to the non-compact setting is only a matter of having appropriate definitions for Floer homology (see again [1] where we follow ideas of Kontsevich and give the relevant example of an appropriate definition of Floer homology of Lagrangians with boundary on a complex hypersurface).
- (3) The correspondence between Morse homology and ordinary homology is due to Morse. The earliest appearance of such an idea in the case of manifolds with boundary is probably [4]. Because of the specific Morse complex considered, the author was unable to find a reference which proved exactly the desired statement at the level of homology, although nothing original needs to be performed in order to achieve this, so it seems appropriate to simply credit Morse and Baiada.
- (4) The fact that one can interpret the Čech cohomology of line bundles on a toric variety in terms of the ordinary homology of the corresponding fan is well known among toric geometers and appears in standard textbooks such as [18]. We are taking the dual point of view by working with a moment polytope.

In particular, it is clear that at the level of homology, every step after the first is a generalization of results that are well established. We now describe, in more detail, how to proceed in the A_∞ setting.

- (0) To formalize the issues of transversality (rather, lack thereof), we use unital A_∞ pre-categories as a technical tool. Conjecturally, see [28], the homotopy category of such objects is equivalent to that of A_∞ categories, although we prove, and use, no result of this type.
- (1) Consider the Laurent polynomial

$$W: (\mathbb{C}^*)^n \rightarrow \mathbb{C}$$

which is conjecturally the mirror of X , and the map

$$\mathrm{Log}: (\mathbb{C}^*)^n \rightarrow \mathbb{R}^n$$

which takes an n -tuple of complex numbers to the n -tuple of logarithms of their respective norms. Using the tropical techniques of Mikhalkin, we can pretend that the amoeba of M , its image under the Log map, is a smooth manifold with boundary which is close to a polyhedral complex Π called the tropical amoeba. It was observed in [1, Section 3.2] that, for appropriate choices of tropical degenerations, the moment polytope Q of X appears as a component of the complement of Π . In particular, there is a smoothing \mathcal{Q} of Q which is a component of the complement of the amoeba of M .

Let $\mathcal{F}\text{uk}((\mathbb{C}^*)^n, M)$ be the sub-pre-category of $\text{Fuk}((\mathbb{C}^*)^n, M)$ whose objects are Lagrangian sections of the restriction of Log to \mathcal{Q} such that their boundary is a subset of M , and which are admissible (see Definition 3.7).

- (2) Since our pre-category $\mathcal{F}\text{uk}((\mathbb{C}^*)^n, M)$ consists of sections of a Lagrangian torus fibrations, we can compute Floer cohomology groups by passing to the universal cover and taking all possible lifts, remembering the grading by $\mathbb{Z}^n = \pi_1(T^n)$. Since the universal cover is essentially the cotangent bundle of the disc, we can use the results of [17] which prove that for compact manifolds without boundary, the moduli spaces of holomorphic discs and of gradient trees agree (exactly!) once the number of marked points is fixed, and the complex structure on the cotangent bundle is appropriately rescaled. It suffices therefore to develop the correct versions of Floer and Morse theory in the case of manifolds with boundary, and to address the fact that Fukaya and Oh do not prove the existence of an almost complex structure which makes the moduli spaces of holomorphic discs and gradient trees agree regardless of the number of marked points (the author knows of no reason for such an almost complex structure to exist). The relevant Morse and Floer theories for manifolds with boundary are developed in this paper, while we use some techniques of homological algebra in Section 4.3 to construct A_∞ quasi-isomorphisms relating the Fukaya pre-category and the Morse pre-category $\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$ despite the lack of an almost complex structure which would make the two pre-categories isomorphic.
- (3) Every smooth function on \mathcal{Q} yields a decomposition of the boundary of \mathcal{Q} into a subset $\partial^+ \mathcal{Q}$ where the function is increasing, and another where it is decreasing. The Morse cohomology of such a function is in fact computing the relative cohomology group

$$H^*(\mathcal{Q}, \partial^+ \mathcal{Q}).$$

Further, we define moduli spaces of trees which allow us to relate the A_∞ structure on Morse functions to the cup product on the cellular chains of an appropriate triangulation. A key property of the functions which appear in $\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$ is that the subset $\partial^+ \mathcal{Q}$ essentially contains a union of cells of the polytope ∂Q as a deformation retract. This means that this cohomology group can also be understood combinatorially. By interpreting the chains of the dual barycentric subdivision of the natural cell decomposition of Q as cochains on Q_b , and studying moduli spaces of gradient trees with boundary conditions on cells of this barycentric subdivision, we construct an A_∞ quasi-isomorphism from a combinatorial pre-category $\text{Simp}^{\mathbb{Z}^n}(Q_b)$ to $\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$.

- (4) We strengthen the classical result of toric geometry by proving that the category $\text{Simp}^{\mathbb{Z}^n}(Q_b)$ is equivalent as a DG category to the category $\check{\text{Cech}}(X)$ of line bundles on X with morphisms given by the Čech complexes and composition given by the cup product.

3. THE CATEGORY OF TROPICAL LAGRANGIAN SECTIONS

3.1. The Fukaya pre-category of a superpotential. Strictly speaking we only defined the Donaldson-Fukaya pre-category of admissible Lagrangians in [1]; i.e. we only studied the cup product at the level of homology. Before proceeding to the case of Landau-Ginzburg potentials, let us recall the definition of the Fukaya pre-category in the absence of a superpotential. We will work with manifolds whose symplectic form is exact which greatly simplifies the technical difficulties; for an overview of the theory, see [40, Chapter 2]. Let N be a complete Weinstein manifold, in other words, N is equipped with both a symplectic form ω and a 1-form λ such that

$$d\lambda = \omega,$$

and the corresponding Liouville flow is a gradient flow for some proper Morse function on N . For details, see [10]. Assume further that the first Chern class of N vanishes. Choose an almost complex structure on N which is convex at infinity, and a complex volume form, i.e. a trivialization of the (complex) canonical bundle. The reader should have in mind the example where N is the cotangent bundle of a Riemannian manifold V , equipped with canonical 1-form pdq and the complex volume form induced by the real volume form on V .

The Fukaya category of N should have, as objects, Lagrangian submanifolds of N . However, even though N is exact, all difficulties do not disappear, so we will restrict to Lagrangians L which satisfy the following conditions:

- The submanifold L is equipped with a spin structure. This assumption is necessary in order to obtain natural orientations on the moduli spaces of holomorphic polygons which will be used in the definition of the Fukaya category. In particular, we can only part with this condition if we are willing to work only in characteristic 2,
- The restriction of λ to L is exact. By a standard argument involving Stokes' theorem, this implies that there are no holomorphic discs with boundary on L , which in turn implies that the differential defined on Floer complexes is an honest differential; i.e. its square vanishes.
- We are given a choice of grading on L . Gradings are necessary in order for the Floer chain complex to be, as one might guess, \mathbb{Z} -graded. Without such a choice, they are only naturally $\mathbb{Z}/2\mathbb{Z}$ graded. A complex volume form on N induces a real structure on its canonical bundle. At every point on L , the tangent space TL determines a real line in the canonical bundle. A grading is a choice of an \mathbb{R} valued function whose exponential gives the S^1 valued angle between this line and the x -axis. Of course, if $H_1(L) = 0$, such a grading always exists. For a thorough discussion of grading issues in Floer theory, see [39].

Definition 3.1. *A Lagrangian brane $L \subset N$ is an embedded compact Lagrangian submanifold equipped with the above choices of additional structures.*

Given transverse Lagrangian branes L_1 and L_2 in N , let $\text{CF}^*(L_1, L_2)$ be the graded free abelian group generated by the intersection points $L_1 \cap L_2$. Again, the reader can refer to [39] for how to assign gradings to intersection points. If L_1 and L_2 are graphs of exact 1-forms in the cotangent bundle of a manifold V of dimension n , an intersection point p corresponds to a critical point of the difference between primitives f_1 and f_2 for the 1-forms defining L_1 and L_2 . In this case,

$$\deg(p) = n - \mu(p),$$

where $\mu(p)$ is the Morse index of p as a critical point of $f_2 - f_1$.

Given an almost complex structure J on N , Floer considered the moduli space of holomorphic maps

$$u: \mathbb{R} \times [0, 1] \rightarrow N$$

which, for all t , take $(t, 0)$ to L_1 and $(t, 1)$ to L_2 . If u has finite energy, he proved that $u(t, s)$ must converge to intersection points p and q between L_1 and L_2 as t converges to $\pm\infty$. We denote the moduli space of such maps by

$$\mathcal{M}(q, p).$$

Translation in the \mathbb{R} factor of the domain defines a free \mathbb{R} action on this space. We write

$$\mathcal{M}(q, p)/\mathbb{R}$$

for the quotient. If L_1 and L_2 are generic (a C^∞ small Hamiltonian perturbation of one of the Lagrangians suffices, see [36]) we can now define an endomorphism of $\text{CF}^*(L_1, L_2)$ by the formula

$$dp = \sum_{q \ni \dim(\mathcal{M}(q, p))=1} |\mathcal{M}(q, p)/\mathbb{R}| \cdot q.$$

In the above formula, $|\mathcal{M}(q, p)/\mathbb{R}|$ is a signed count of the number of components of $\mathcal{M}(q, p)$. Since we have discussed neither orientations on $\mathcal{M}(q, p)$ nor gradings on the Floer complex, we refer to [40] for the sign considerations which are needed to prove of the following Lemma. Without signs, this result is due to Floer [12], and is by now a standard application of Gromov's compactness theorem [21].

Lemma 3.2. *The map d is a differential on $\text{CF}^*(L_1, L_2)$.* □

One would like to define compositions

$$\text{CF}^*(L_2, L_3) \otimes \text{CF}^*(L_1, L_2) \rightarrow \text{CF}^*(L_1, L_3)$$

by using moduli spaces of curves with boundary on L_1 , L_2 and L_3 . The resulting composition will, however, only be associative after passing to homology. Fukaya observed that even at the chain level, one can construct a strongly homotopy associative composition by considering all moduli spaces of holomorphic polygons. We briefly describe his construction:

Let (L_0, \dots, L_n) be a sequence of Lagrangian branes in general position, and let S be a collections of points $(s_{0,1}, s_{1,2}, \dots, s_{d-1,d}, s_{0,d})$ on the boundary of the closed disc D^2 ordered counter-clockwise. We consider finite energy holomorphic maps

$$u: D^2 - S \rightarrow N$$

such that the image of the interval from $s_{i-1,i}$ to $s_{i,i+1}$ lies on L_i . The finite energy condition implies that u extends to a continuous map on D^2 taking $s_{i,i+1}$ to an intersection point $p_{i,i+1}$ between L_i and L_{i+1} . We define

$$\mathcal{M}(p_{0,d}; p_{0,1}, \dots, p_{d,d-1})$$

to be the moduli space of such maps. Again, the choices of spin structures on L_i induce an orientation on this moduli space. Further, assuming the Lagrangians are in generic position, these moduli spaces are smooth manifolds of dimension

$$(3.1) \quad d - 2 + \deg(p_{0,d}) - \sum_i \deg(p_{i,i+1}).$$

The genericity condition used here is the assumption that moduli spaces of holomorphic polygons are regular, see [13].

Fukaya defined an operation μ_d of degree $2 - d$

$$\mu_d: \text{CF}^*(L_{d-1}, L_d) \otimes \dots \otimes \text{CF}^*(L_0, L_1) \rightarrow \text{CF}^*(L_0, L_d).$$

On generators, μ_d is given by the formula

$$(3.2) \quad p_{d-1,d} \otimes \dots \otimes p_{0,1} \mapsto \sum_{\deg(p_{0,d})=2-d+\sum_i \deg(p_{i,i+1})} |\mathcal{M}(p_{0,d}; p_{0,1}, \dots, p_{d,d-1})| p_{0,d},$$

where $|\mathcal{M}(p_{0,d}; p_{0,1}, \dots, p_{d,d-1})|$ is the signed number of components of the zero-dimensional manifold $\mathcal{M}(p_{0,d}; p_{0,1}, \dots, p_{d,d-1})$. Gromov's compactness theorem implies the following result, whose proof with signs is given in [40, Chapter 2].

Lemma 3.3. *The operations μ_d satisfy the A_∞ equation. Explicitly, we have*

$$(3.3) \quad \sum_{i+d_2 \leq d} (-1)^{\mathfrak{X}_i} \mu_{d-d_2+1}(\mathbb{1}^{d-d_2-i} \otimes \mu_{d_2} \otimes \mathbb{1}^i) = 0,$$

where \mathfrak{X}_i depends on the degree of the first i inputs as follows:

$$\mathfrak{X}_i = i + \sum_{j=0}^{i-1} \deg(p_{j,j+1}).$$

□

We can summarize the above discussion in the following definition. We refer the reader to Appendix D for the definition of A_∞ pre-categories.

Definition 3.4. *The **Fukaya pre-category** of a complete Weinstein manifold N is the A_∞ pre-category $\text{Fuk}(N; J)$ given by the following data.*

- *Objects are Lagrangian branes in N .*
- *A sequence (L_0, \dots, L_d) is transverse whenever all moduli spaces of holomorphic polygons with boundary on a subsequence are regular.*
- *The space of morphisms between a transverse pair (L_0, L_1) is the Floer complex $\text{CF}^*(L_0, L_1)$.*
- *The higher compositions are given by the maps μ_d .*

Remark 3.5. We will often drop J from the notation, and simply write $\text{Fuk}(N)$

Floer's original applications of Floer homology relied on its invariance under Hamiltonian isotopies. Using the composition μ_2 , one can state a stronger result

as follows. If L'_0 is a Hamiltonian deformation of L_0 , then there exists a cycle $e \in \text{CF}^0(L_0, L'_0)$ such that

$$\mu_2(e, -): \text{CF}^*(L_1, L_0) \rightarrow \text{CF}^*(L_1, L'_0)$$

induces an isomorphism on cohomology whenever (L_1, L_0, L'_0) is a transverse triple in the Fukaya pre-category. Reversing the roles of L_0 and L'_0 , we obtain the following result.

Lemma 3.6. *The pre-category $\text{Fuk}(N)$ is unital.* \square

3.1.1. The Fukaya Category of a Landau-Ginzburg potential. Consider a map $f : N \rightarrow \mathbb{C}$ which is a symplectic fibration away from its critical values and assume that the zero fibre is a smooth manifold M . Assume further that we have chosen an almost complex structure on N which respect to which f is holomorphic in a neighbourhood of M . The following definition appears in [1, Section 2], and is a reformulation of a definition of Kontsevich in [27].

Definition 3.7. *An admissible Lagrangian brane is a compact Lagrangian submanifold $\mathcal{L} \subset N$ with boundary on M together with a choice of grading and spin structure, satisfying the following conditions:*

- *The restriction of λ to \mathcal{L} is exact, and*
- *There exists a small neighbourhood of $\partial\mathcal{L}$ in \mathcal{L} which agrees with the parallel transport of $\partial\mathcal{L}$ along a segment in \mathbb{C} .*

Note that the last two conditions are vacuous if $\partial\mathcal{L}$ is empty, which is a possibility that we do not exclude.

Next, we consider the conditions under which the higher products are well behaved.

Definition 3.8. *A sequence of admissible Lagrangians $(\mathcal{L}_0, \dots, \mathcal{L}_n)$ is **positive** if the curves γ_k which agree with $f(\mathcal{L}_k)$ near the origin all lie in the left half plane and their tangent vectors v_k are oriented counter-clockwise (See Figure 3).*

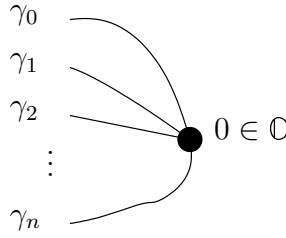


FIGURE 3. Curves corresponding to positively oriented Lagrangians.

Note that if $(\mathcal{L}_i)_{i=1}^n$ is a positive sequence then the angle between v_i and v_j is again positive whenever $i < j$. In particular, every subsequence of a positive sequence is positive.

Let us assume that the Lagrangians \mathcal{L}_i intersect transversally in the interior. Throughout, we will only work with interior intersection points. If $p_i \in \mathcal{L}_{i-1} \cap \mathcal{L}_i$ are such interior intersection points between successive Lagrangians, and $p_0 \in \mathcal{L}_d \cap \mathcal{L}_0$, we can define

$$\mathcal{M}(p_0; p_1, \dots, p_d)$$

to be the moduli space of holomorphic polygons with boundary on $\bigcup_i \mathcal{L}_i$ and corners at p_i . If we perturb each Lagrangian \mathcal{L}_i near the intersections by a C^∞ small amount, we can guarantee that all moduli space of holomorphic polygons are regular, see [13]. This allows us to define the Gromov bordification

$$\overline{\mathcal{M}}(p_0; p_1, \dots, p_d)$$

by including “cusp-curves”.

Lemma 3.9. *If $\{\mathcal{L}_i\}$ is positive, then the Gromov bordification is compact.*

Proof (c.f. Lemma A.13). It suffices to prove that curves in

$$\mathcal{M}(p_0; p_1, \dots, p_d)$$

cannot escape to the boundary. The proof given in [1, Lemma 2.8] can be used here. The key idea is to project by f to \mathbb{C} , and observe that if u_s is a family of holomorphic maps parameterized by $[0, 1]$ such that the image of u_0 does not intersect a neighbourhood of the origin, then the same property holds for u_1 . \square

If we define the higher products μ_d by counting the number of isolated holomorphic polygons as in the case where the Lagrangians have no boundary, we conclude

Corollary 3.10. *If all moduli spaces are regular, then the A_∞ equation holds.* \square

We are now ready to define the main object which we shall be studying.

Definition 3.11. *The **Fukaya pre-category** of the Landau-Ginzburg potential f is the A_∞ pre-category $\text{Fuk}(N, M)$ given by the following data:*

- *Objects are admissible Lagrangian branes.*
- *Transverse sequences are positive sequences of Lagrangian branes such that all holomorphic polygons with boundary values on a sub-sequence are regular.*
- *Morphisms are the Floer complexes with the differential counting holomorphic bi-gons and higher products counting holomorphic polygons.*

3.2. Tropical Lagrangians for mirrors of toric varieties. Let X be a smooth toric variety defined by a fan Δ . We assume X to be projective, so $|\Delta| = \mathbb{R}^n$ and there exists an integral piece-wise linear function

$$\psi : \mathbb{R}^n \rightarrow \mathbb{R}$$

which is linear on each cone of Δ , and which is strictly convex. Let Q be the moment polytope of the line bundle which corresponds to ψ . In [1, Definition 4.3], we constructed a cover $\{O_\tau\}$ of \mathbb{R}^n labeled by the faces of Q such that the complement in τ of a small neighbourhood of $\partial\tau$ is contained in O_τ .

Remark 3.12. We are about to extensively use results from [1]. The notation has been kept relatively consistent, with one major difference. In [1], we used the greek letters σ and τ to label cones in the fan Δ . In this paper, we use the same letter to label the faces of the polytope Q . The open sets O_τ therefore correspond to $O_{\tilde{\tau}}$ in the notation of [1].

We observed that the choice of ψ induces a tropical degeneration of the mirror of X . In particular, we have a family of Laurent polynomials W_t , whose zero fibre are (for $t \gg 0$) are smooth complex (and hence symplectic) hypersurfaces which we denote by $M_t \subset (\mathbb{C}^*)^n$. We considered the amoebas \mathcal{A}_t of M_t , i.e the

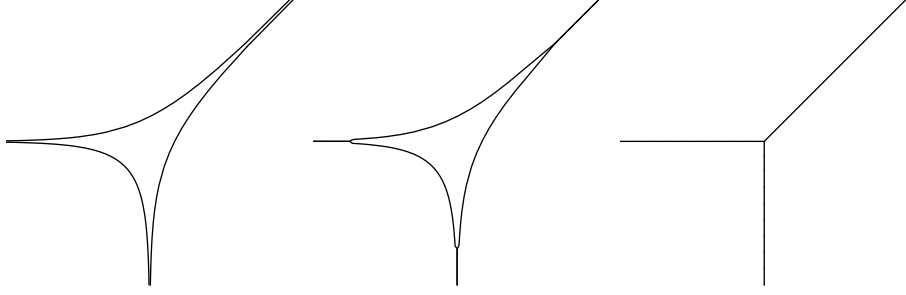


FIGURE 4. The amoeba of the standard hyperplane in $(\mathbb{C}^*)^2$, its “tropical localization,” and its tropical amoeba

projection of M_t to \mathbb{R}^n under the logarithmic moment map. Further, we constructed a “tropical localization of M_t ” by deforming W_t to a (non-algebraic) map $W_{t,1}$ so that the resulting hypersurface $M_{t,1}$ is particularly well described by its amoeba. To understand the relation between the amoebas of M , $M_{t,1}$, and the tropical amoeba, it is best to consider the case of the standard hyperplane in $(\mathbb{C}^*)^2$ which is illustrated in Figure 4.

By construction, the tangent space of any cell τ is an integral subspace of \mathbb{R}^n , i.e. it is defined as the vanishing locus of finitely many integral covectors. In particular, if $T_{\mathbb{Z}}\tau$ is the intersection of this tangent space with the lattice \mathbb{Z}^n , then we have an identification

$$T\tau = T_{\mathbb{Z}}\tau \otimes_{\mathbb{Z}} \mathbb{R}.$$

Let us write $\tilde{\tau}$ for the linear span of the covectors which vanish on $T\tau$, and $(\mathbb{C}^*)^{\tilde{\tau}}$ for the corresponding complex torus. The natural map

$$(\mathbb{C}^*)^n \rightarrow (\mathbb{C}^*)^{\tilde{\tau}}$$

has fibre which can be identified with $(\mathbb{C}^*)^{\tau} = T\ker(\tilde{\tau})/T_{\mathbb{Z}}\ker(\tilde{\tau})$ (the quotient of the tangent space of the linear plane tangent to τ by the lattice). Note that the choice of a splitting of the inclusion $T_{\mathbb{Z}}\tau \hookrightarrow \mathbb{Z}^n$, gives a product decomposition

$$(\mathbb{C}^*)^n \cong (\mathbb{C}^*)^{\tilde{\tau}} \times (\mathbb{C}^*)^{\tau}$$

of complex manifolds. We can now restate, in a slightly enhanced form, a results which appeared in [1, Lemma 4.4]:

Lemma 3.13. *The intersection of the hypersurface $M_{t,1}$ with the inverse image of $\log(t) \cdot O_{\tau}$ in $(\mathbb{C}^*)^n$, splits as a product*

$$M_{t,1}^{\tilde{\tau}} \times (\tau \cap O_{\tau}) \times \mathbb{T}^{\tau}$$

where

- $M_{t,1}^{\tilde{\tau}}$ is the tropical localization of a hypersurface in a torus $(\mathbb{C}^*)^{\tilde{\tau}}$, and
- $\mathbb{T}^{\tau} = T\tau/T_{\mathbb{Z}}\tau$ is the real subtorus of $(\mathbb{C}^*)^{\tau}$.

□

Unfortunately, this is not a product decomposition as symplectic manifolds. In order to obtain such a decomposition, we would essentially have to pick the splitting of the inclusion $T_{\mathbb{Z}}\tau \hookrightarrow \mathbb{Z}^n$ compatibly with the euclidean metric, which is of course not always possible. In general, the orthogonal complement of $T_{\mathbb{Z}}\tau$ in \mathbb{Z}^n projects

to a finite index sub-lattice of the quotient $\mathbb{Z}^n/T_{\mathbb{Z}}\tau$. If we write $(\mathbb{C}^\star)^{\tau^\perp}$ for the corresponding torus, we obtain a finite cover

$$(\mathbb{C}^\star)^\tau \times (\mathbb{C}^\star)^{\tau^\perp} \rightarrow (\mathbb{C}^\star)^n.$$

The inverse image of $M_{t,1}$ in the cover splits as a product in both the symplectic and complex categories. The factor living in $(\mathbb{C}^\star)^{\tau^\perp}$ now corresponds to a higher degree hypersurface (rather than a hyperplane).

In [1, Lemma 4.16] we observed that one of the components $\mathcal{Q}_{t,1}$ of the complement of $\mathcal{A}_{t,1}$ satisfies the following conditions:

- $\mathcal{Q}_{t,1}$ is contained in (and is C^0 close to) the $\log(t)$ rescaling of the moment polytope Q , and
- the boundary of $\mathcal{Q}_{t,1}$ is smooth, and its image under the zero section lies in M .

Note that the boundary $\mathcal{Q}_{t,1}$ again decomposes as a product in O_τ . Since we are now working in \mathbb{R}^n rather than in the complex torus, we can choose this product decomposition to be compatible with the metric. More precisely:

Lemma 3.14. *The restriction of $\mathcal{Q}_{t,1}$ to $\log(t) \cdot O_\tau$ splits as a Riemannian product*

$$\mathcal{Q}_{t,1}^\perp \times (\tau \cap O_\tau),$$

where $\mathcal{Q}_{t,1}^\perp$ is the amoeba of the tropicalisation of a hyperplane. □

Remark 3.15. In order to see $\mathcal{Q}_{t,1}^\perp$ as the amoeba of a hyperplane (rather than a higher degree hypersurface) we equip the plane orthogonal to τ with the lattice coming from its identification with the quotient of \mathbb{R}^n by the tangent space to τ .

From now on (except for Section 5.2 and Appendix A), we shall drop the subscripts from $M_{t,1}$, $W_{t,1}$, and the associated objects. We will further rescale \mathbb{R}^n by $\frac{1}{\log(t)}$, and write \mathcal{Q} for the image of $\mathcal{Q}_{t,1}$ under this rescaling. Note that \mathcal{Q} is C^0 close to the original polytope Q .

We denoted the image of \mathcal{Q} under the zero section by \mathcal{L} . It is supposed to correspond under mirror symmetry to the structure sheaf of X .

Definition 3.16. *The pre-category $\mathcal{T}\text{Fuk}((\mathbb{C}^\star)^n, M)$ of **tropical Lagrangian sections**, is the A_∞ pre-category of admissible Lagrangians branes in $((\mathbb{C}^\star)^n, M)$ which are sections of the restriction of the moment map to \mathcal{Q} .*

Remark 3.17. It would be more consistent to use the terminology “semi-tropical Lagrangian section,” hinting that we did not take the tropical degeneration to its limit, but rather stopped just before the limit point, with the smooth symplectic structure still present, yet the tropical features apparent.

Given \mathcal{L}_1 and \mathcal{L}_2 two tropical Lagrangian sections, the space of morphisms $\text{CF}^i(\mathcal{L}_1, \mathcal{L}_2)$ is naturally graded by an affine space over \mathbb{Z}^n . This is a grading which is independent of the cohomological grading. Indeed, since we have fixed a zero section \mathcal{L} and a fixed lift $\tilde{\mathcal{L}}$ to the universal cover, we can consider the difference $\mathcal{L}_2 - \mathcal{L}_1$, which yields a new section of the fibration. The lifts of $\mathcal{L}_2 - \mathcal{L}_1$ to the universal cover form an affine space over \mathbb{Z}^n , which we identify with the grading of $\text{CF}^i(\mathcal{L}_1, \mathcal{L}_2)$. Given an intersection point p between \mathcal{L}_2 and \mathcal{L}_1 , we obtain an intersection point between $\mathcal{L}_2 - \mathcal{L}_1$ and the zero section, and the corresponding lift

to the universal cover is uniquely determined by the fact that \tilde{p} , the lift of p , lies on $\tilde{\mathcal{L}}$.

A priori, there are uncountably many tropical Lagrangian sections. As we shall see, there are in fact only countably many Hamiltonian isotopy classes, naturally in bijective correspondence with isomorphism classes of line bundles on X . Recall that

$$\text{Log} : (\mathbb{C}^*)^n \rightarrow \mathbb{R}^n$$

is the moment map taking an n -tuple of non-zero complex numbers to the logarithms of their absolute values. We abusively write $\text{Log}|_{\mathcal{Q}}$ for the restriction of this torus fibration to the inverse image of the disc \mathcal{Q} .

Lemma 3.18. *Every section of $\text{Log}|_{\mathcal{Q}}$ which corresponds to a tropical Lagrangian section agrees with the zero section when restricted to*

$$\partial\mathcal{Q} \cap O_v,$$

whenever v is a vertex of Q .

Proof. It is easy to check that the boundary of the amoeba of the standard hyperplane $\sum_{i=1}^n z_i = 1$ agrees with its real locus. In particular, the restriction of the Logarithmic moment map to the standard hyperplane in $(\mathbb{C}^*)^n$ is injective when restricted to the inverse image of the boundary of its amoeba. The results of [1, Lemma 4.4] restated above imply that the amoeba of M is modeled, in a neighbourhood of every vertex of Q , after the amoeba of the standard hyperplane. Since the zero section is a tropical Lagrangian, every tropical Lagrangian section must therefore agree with it on this subset of the boundary. \square

In particular, any lift of a tropical Lagrangian section to $T^*\mathbb{R}^n$ must take 2π integral values. As an extension of the previous lemma, we have

Lemma 3.19. *The lift of any tropical Lagrangian section when restricted to*

$$\partial\mathcal{Q} \cap O_\tau,$$

lies in an affine subspace of \mathbb{R}^n which is a 2π -integral translate of the tangent space of τ . A Lagrangian section which satisfies this property for every τ is isotopic through Lagrangian sections with boundary on $M_{t,1}$ to a tropical Lagrangian section.

Proof. We first prove that any section (not necessarily Lagrangian) of the restriction of $T^*\mathbb{R}^n$ to $\partial\mathcal{Q} \cap O_\tau$ which lands in M must have values in a 2π -integral translate of the tangent space of τ . The previous lemma proves the base case when τ is a vertex and the induction step follows easily from the product decomposition of Lemma 3.13.

To prove the second part, we must achieve admissibility. Consider a function f defining a Lagrangian section Λ_f with boundary $\partial\Lambda_f \subset M_{t,1}$. Using the parallel transport maps defined by the symplectic connection induced by W , we can construct a Lagrangian Λ_g whose boundary agrees with that of Λ_f . The arguments of [1, Section 5] imply that Λ_g (near $\partial\mathcal{Q}_{t,1}$) is defined by a function g . By linearly interpolating between f and g , we obtain the desired isotopy. \square

We are now ready to prove

Proposition 3.20. *The Hamiltonian isotopy class of a tropical Lagrangian section is determined by the values of its lifts near the vertices of Q*

Proof. Every such lift is the graph of the differential of a function on the disc. Let f_1 and f_0 denote two such functions whose differentials are equal near the vertices. Consider the 1-parameter family

$$f_t = tf_0 + (1-t)f_1.$$

From the characterization of the previous lemma, it easily follows that the Lagrangians $\Lambda_{f_t} \subset (\mathbb{C}^*)^n$ defined by f_t have boundary on M . In particular, there is a family g_t of functions whose graphs define admissible Lagrangians with boundaries $\partial\Lambda_{f_t}$. Picking such a family with $f_0 = g_0$ and $f_1 = g_1$ yields the desired isotopy. \square

Corollary 3.21. *There is a bijective correspondence between Hamiltonian isotopy classes of tropical Lagrangian sections of $((\mathbb{C}^*)^n, M)$ and isomorphism classes of line bundles on X .*

Proof. It is well known that isomorphism classes of line bundles are given by piecewise linear integral function on \mathbb{R}^n , with domains of linearity the maximal cones of Δ , modulo the (global) addition of an integral linear function on \mathbb{R}^n . The data of a lattice vector for each vertex of Q corresponds to a linear integral function on the maximal cones of Δ . The fact that the restriction of our section to an edge of Q can only vary by multiples of the integral tangent vector of the edge implies that the two linear functions associated to the maximal cones dual to the endpoints of an edge agree on the hypersurface where they meet. Starting with a piecewise linear integral function on Δ , it is easy to construct the corresponding section by a patching argument. By Lemma 3.19 this Lagrangian with boundary on M is isotopic to an admissible one, hence to a tropical Lagrangian section. Finally, we note that addition of an integral linear function on \mathbb{R}^n corresponds to choosing different lifts. \square

In Sections 5 and 6 we will prove that this bijective correspondence extends to an equivalence of A_∞ pre-categories. In order to do this, we shall have to develop some general machinery relating the multiplicative structures on the chain complexes defining Morse, Floer, Čech, cellular, and simplicial cohomology theories.

4. GENERAL FRAMEWORK

In this section, we construct chain-level A_∞ quasi-equivalences between some pre-categories where the spaces of morphism are respectively Morse, Floer, cellular, and Čech complexes. We will work with integral coefficients throughout this section since no serious difficulties are simplified by working over a field. First, we introduce the relevant Morse theory.

4.1. Morse theory on manifolds with boundary. Let $(Q, \partial Q)$ be a smooth oriented closed Riemannian manifold. It is well known that one can use Morse theory to study the topology of the pair $(Q, \partial Q)$ without any properness condition. The reference which closest in spirit to the approach which we need is [25]. However, the exact results that we'll use do not seem to be written down anywhere, so we give a brief description.

Definition 4.1. *A function $f: Q \rightarrow \mathbb{R}$ is **boundary convex** if no singular gradient trajectory between interior points intersects the boundary.*

Remark 4.2. Recall that if f is not Morse, then a singular gradient trajectory is a concatenation of gradient trajectories and paths contained in the critical locus.

There is an important local model which justifies our terminology for boundary convex functions:

Lemma 4.3. *A subset C of \mathbb{R}^n has convex boundary if and only if the restriction of every linear function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ to C is boundary convex.* \square

In general, the boundary convexity of a function is a global property. However, we will exploit the existence of sufficient local condition:

Lemma 4.4. *Let $f: Q \rightarrow \mathbb{R}$ be a function whose interior critical points are isolated from the boundary. Then f is boundary convex if every point $r \in \partial Q$ at which $\text{grad}(f)$ is not transverse to the boundary satisfies one of the following two conditions:*

- (1) *If r is a critical point of f and γ is a path starting at r which is not tangent to the boundary then $f|_\gamma$ is strictly decreasing in some neighbourhood of r .*
- (2) *If r is not a critical point of f , then (if it is defined) a gradient segment starting at r is contained in ∂Q .*

Proof. We assume, by contradiction, that f satisfies the above property, and that there are interior points p and q which are connected by a singular gradient trajectory γ which is tangent to the boundary at some point. Let r be the last point at which γ intersects ∂Q . Since f is non-decreasing along γ , the two conditions above preclude the existence of such a point, hence the existence of γ . \square

A function satisfying the conditions of the above Lemma will be called **positively boundary convex**.

4.1.1. *The Morse complex.* The material in this section is familiar in the proper case, see for example [38], [15], and [5]. We will restate some definitions in a slightly non-standard way. Recall that for each non-degenerate critical point p of a smooth function f we have an ascending (stable) manifold

$$W^s(p)$$

consisting of points whose image under the negative gradient flow converges to p as $t \rightarrow +\infty$, and a descending (unstable) manifold

$$W^u(p)$$

consisting of points whose image under the positive gradient flow converges to p as $t \rightarrow +\infty$.

Definition 4.5. *The degree of p , denoted $\deg(p)$, is the dimension of $W^s(p)$.*

Even though Q has been assumed to be oriented, the submanifolds $W^s(p)$ do not inherit a natural orientation.

Definition 4.6. *The orientation line at p is the abelian group generated by the two orientations of $W^s(p)$, with the relation that the two opposite orientations satisfy*

$$[\Omega_p] + [-\Omega_p] = 0.$$

We denote this group by $|o_p|$.

Definition 4.7. *If f is a boundary convex function on a smooth manifold Q , we define $\partial_f^+ Q$ to be the set of points in ∂Q consisting of limit points of gradient trajectories of f with initial points in the interior of Q . The complement of $\partial_f^+ Q$ in ∂Q is denoted $\partial_f^- Q$.*

We find the following definition convenient, even though it may not agree with some conventions which require Morse-Smale functions to be proper.

Definition 4.8. *A function with non-degenerate critical points is called **Morse-Smale** if all ascending and descending submanifolds of interior critical points intersect each other transversally, and $\partial_f^+ Q$ is a compact codimension 0 submanifold of the boundary.*

Given a boundary convex Morse-Smale function f , we can now define the Morse complex

$$\mathrm{CM}^i(f) = \bigoplus_{\substack{p \in Q - \partial Q \\ \deg(p)=i}} |o_p|.$$

To define the differential, we begin by observing that since Q is oriented, an orientation of any ascending manifold induces an orientation of the corresponding descending manifold. Further, if

$$\deg(q) = \deg(p) + 1,$$

then $W^u(p)$ and $W^s(q)$ intersect in a 1-dimensional submanifold of Q , consisting of finitely many components whose image under f is a fixed open segment in \mathbb{R} . So if we pick orientations on the ascending manifolds at p and q , we can compare the induced orientation on

$$W^u(p) \cap W^s(q)$$

with the natural orientation pulled back from \mathbb{R} using f . Let $n(q, p)$ denote the difference between the number of components on which the two orientations agrees, and the ones on which they differ. This number depends on the choice of orientations, but only up to sign. We can now define the differential

$$\mu_1: \mathrm{CM}^i(f) \rightarrow \mathrm{CM}^{i+1}(f)$$

by choosing orientations Ω_q of each ascending manifold and linearly extending the formula

$$\Omega_p \mapsto \sum_q n(q, p) \Omega_q.$$

Note that even though we chose orientations to write a formula, the differential is independent of all choices. Let us also note that we are counting gradient lines for $-f$ starting at p and ending at q . This causes some unfortunate sign conventions in the next section.

In general, the differential on the Morse complex of a Morse-Smale function on a manifold with boundary may not square to 0. However, we have

Lemma 4.9. *If f is a boundary convex Morse-Smale function, then $\mu_1^2 = 0$.*

Proof. The same proof usually given in the compact case works. The moduli space of gradient trajectories converging at either ends to critical points whose degree differs by 2 is a manifold whose boundary consists precisely of the terms of μ_1^2 . Convexity prevents 1-parameter families of gradient lines from escaping to the

boundary, so this manifold is compact, and the signed count of these boundary points vanishes. \square

4.1.2. The A_∞ structure. The A_∞ structure on the pre-category of Morse functions was discovered by Fukaya, see [14]. Let $\{f_i\}_{i=0}^n$ be a sequence of functions on Q , and assume that each pair $f_i - f_j$ for $i < j$ is boundary convex. We will construct higher products by considering moduli spaces of gradient trees. For concreteness, and to set the stage for Section 4.2, we review their construction.

Definition 4.10. *A singular metric ribbon tree T consists of the following data:*

- (1) *one dimensional finite CW complex,*
- (2) *a map l from the set of edges to the interval $(0, +\infty]$ (the length),*
- (3) *a cyclic ordering of the edges adjacent to each vertex,*
- (4) *a marked univalent vertex called the outgoing vertex, and*
- (5) *a labeling of all vertices as either of finite or infinite type.*

The above data must satisfy the following conditions

- (1) *There are no non-trivial cycles,*
- (2) *all vertices of valence greater than 2 are of finite type,*
- (3) *all vertices of valence 2 are of infinite type,*
- (4) *an edge has infinite length if and only if it has an endpoint of infinite type.*

Univalent vertices of T and the edges which include them are called **external** with all other vertices and edges called **internal**. Note that external edges may have either finite or infinite length.

Remark 4.11. Upon choosing the outgoing vertex, the remaining external vertices (or edges) of T acquire a linear ordering. We will think of them as **incoming** vertices (similarly for incoming edges). We equip the edges of a ribbon tree with a natural orientation as follows:

- The orientation on the incoming external edges points towards the interior of the tree. The outgoing external edge is oriented away from the interior.
- The orientations on all but one of the edges adjacent to a fixed vertex v point towards v .

One can easily prove by induction that an orientation satisfying the above properties exists and is unique. We shall therefore freely speak of the (unique) “outgoing edge at a vertex” (the one oriented away from v) or of the incoming vertices of a subtree. In addition, the internal vertex of a tree which lies on the outgoing edge will be a convenient basepoint for many of our argument. We will call it the **node**.

Definition 4.12. *A singular ribbon tree is **reducible** if it contains a bivalent vertex. It is otherwise called **irreducible**, or, more simply, a **ribbon tree**.*

*If T is a singular ribbon tree then a subtree $S \subset T$ is called an **irreducible component** if it is maximal among irreducible subtrees.*

Lemma 4.13. *Every singular ribbon tree admits a unique decomposition into irreducible components.* \square

In particular, one should think of singular ribbon trees as obtained by taking two ribbon trees with infinite external edges and identifying the univalent vertices of these two edges.

If we remove the univalent vertices of a ribbon tree, we obtain an honest metric space. If a vertex lies on an infinite edge, the labeling of vertices as either of finite or infinite type determines whether the edge is isometric to $[0, +\infty)$ or $(-\infty, +\infty)$, with 0 corresponding to the vertex of finite type.

In particular there is a natural notion of (isometric) automorphisms. In analogy with the Deligne-Mumford compactification, we have

Definition 4.14. *A ribbon tree is **stable** if its group of automorphisms as a metric tree is finite. A singular ribbon tree is stable if so are all its irreducible components.*

Lemma 4.15. *All ribbon trees which are not stable are isometric to $(-\infty, +\infty)$. In particular, a singular ribbon tree is stable unless it includes an edge with two bivalent endpoints.* \square

Definition 4.16. *The d -th **Stasheff moduli space**, \mathcal{T}_d , is the moduli space of ribbon trees with $d+1$ external edges all of which have infinite type.*

The topology of \mathcal{T}_d is such that the operation of contracting the length of an internal edge to zero is continuous. In [41], Stasheff showed that this moduli space is homeomorphic to a ball (see also Fukaya and Oh in [17]).

Lemma 4.17 (Definition-Lemma). *There exists a compactification of \mathcal{T}_d into a compact manifold with boundary $\overline{\mathcal{T}}_d$ called the **Stasheff polyhedron** such that*

$$\partial \overline{\mathcal{T}}_d = \bigcup_{d_1+d_2=d+1} \bigcup_{1 \leq i \leq d} \mathcal{T}_{d_1} \times \mathcal{T}_{d_2}.$$

Moreover, $\overline{\mathcal{T}}_d$ is naturally stratified by smooth manifolds with corners. \square

The compactification is simply obtained by adding stable singular ribbon trees all of whose external vertices again have infinite type, and the topology is such that a sequence of trees with fixed topology having an edge whose length becomes unbounded converges to a singular tree with a bivalent vertex precisely at that edge.

Given a fixed tree $T \in \mathcal{T}_n$ and a sequence $\vec{f} = (f_0, \dots, f_n)$ of functions on Q , we label the incoming edges with the functions $\{f_i - f_{i-1}\}_{i=1}^n$ and the outgoing edge with $f_n - f_0$.

Lemma 4.18. *There exists a unique labeling of the edges of T by functions $f_e: Q \rightarrow \mathbb{R}$, extending the above labeling of the external edges, such that for every vertex v , the sum of the functions labeling incoming edges is equal to the function labeling the (unique) outgoing vertex. Further, every function f_e is of the form $f_i - f_j$ with $j < i$.*

Proof following [17]. From the data of a cyclic orientation at each vertex we can construct a unique (up to isotopy) embedding ι of T in the unit disc with external vertices lying on the unit circle. We can label the components of the complement of $\iota(T)$ by functions f_i such that each external edge is labeled by the difference between the functions associated to the two components which it separates. To get the correct sign, we assign a positive sign to the function which lies to the left (according to the orientation of the edge). Note that we can extend this labeling to the internal edges as well, and that it satisfies the desired balancing conditions at every vertex. \square

The higher algebraic structure of Morse functions is obtained from the topological properties of the space of maps from trees to Q .

Definition 4.19. *Given a sequence of functions \vec{f} , and critical points $p_{i,i+1}$ of $f_{i+1} - f_i$ and $p_{0,d}$ of $f_d - f_0$, the **moduli space of gradient trees***

$$\mathcal{T}(p_{0,d}; p_{0,1}, p_{1,2}, \dots, p_{d-1,d}) = \mathcal{T}(p_{0,d}; \vec{p})$$

is the space of maps $\phi : S \rightarrow Q$ with S an arbitrary element of \mathcal{T}_d , which satisfy the following conditions:

- *The image of each edge e is a gradient line for f_e .*
- *The orientation of e is given by the negative gradient of f_e . Further, ϕ identifies the length parametrization of each edge e with the gradient vector field of f_e .*
- *If e is an external edge, and $f_e = f_i - f_j$, then the image of ϕ converges to $p_{i,j}$ along e .*

Let us observe that this moduli space depends not only on the functions f_i , but also on the Riemannian metric of Q . We will make this dependence explicit as needed. For simplicity, we will also write

$$\mathcal{T}(\vec{f}) = \coprod_{p_{0,d}, \vec{p}} \mathcal{T}(p_{0,d}; \vec{p}).$$

Note that the moduli space of gradient trees is not compact as the length of an internal edge can go to infinity. However, it admits a bordification

$$\overline{\mathcal{T}}(\vec{f})$$

by observing that, if such a length becomes infinite, the gradient tree converges to a union of gradient trees which share common vertices, i.e. the domain tree converges to a singular ribbon tree equipped with a map to Q . In [15], Fukaya proves that $\overline{\mathcal{T}}(\vec{f})$ is compact whenever Q is compact. We have

Lemma 4.20. *If each pair $f_i - f_j$ for $j < i$ is boundary convex, then $\overline{\mathcal{T}}(\vec{f})$ is compact.*

Proof. Let ϕ_t be such a 1-parameter family converging to ϕ_1 , such that there is a point $x \in T$ for which

$$\phi_1(x) \in \partial Q.$$

If x lies in the interior of an edge e , then the image of every point on e converges to the boundary by the boundary convexity of the function $f_i - f_j$. Since the external edges cannot converge to the boundary, it suffices to prove the following claim.

Claim. If v is a vertex of T and $\phi_1(v) \in \partial Q$, then every edge adjacent to v lies in ∂Q

Proof of Claim. Assume otherwise. To simplify the argument, we re-orient the outgoing edge at v so that it now points towards v , and take the negative of the corresponding function. We now have

$$\sum_{v \in e} f_e = 0.$$

However, if ν is the outwards pointing normal vector to the boundary at $\phi_1(v)$ then

$$\langle \nu, \text{grad}(f_e) \rangle \leq 0,$$

since tracing back the gradient line of f_e pushes us towards the interior of the manifold. Comparing the two equations, we conclude that these inner products must all vanish, i.e. that all gradient lines are tangent to the boundary at v . By boundary convexity, the corresponding edges are contained in ∂Q . \square

The same argument which yields the result when the boundary is empty therefore applies in the boundary convex case to complete the proof of Lemma 4.20. \square

To extract algebraic invariants, we will assume genericity of the functions f_i .

Definition 4.21. *A sequence of functions \vec{f} is **Morse-Smale** if all functions $\{f_i - f_j\}_{i < j}$ are Morse-Smale, and all finite collections of their ascending and descending manifolds are in “general position”. Such a collection is **boundary convex** if all functions $\{f_i - f_j\}_{i < j}$ are boundary convex.*

The precise meaning of general position is a requirement on the regularity of all moduli spaces of gradient tree as explained in Appendix B, in particular Definition B.3. As before, an orientation of the ascending manifolds of the critical points of $f_{i+1} - f_i$ and of $f_d - f_0$ induces an orientation of $\overline{\mathcal{T}}(\vec{f})$; the explanation is again relegated to Appendix C. In particular, transversality implies that if the degrees of the various critical points satisfy

$$(4.1) \quad 2 - d + \sum_i \deg(p_{i,i+1}) = \deg(p_{0,d}),$$

then the corresponding moduli space of gradient trees consists of (signed) points, so we obtain integers

$$n(p_{0,d}, p_{d-1,d}, \dots, p_{0,1})$$

by choosing orientations of the ascending manifolds, and counting the (signed) number of such points. We define the above integers to be 0 whenever Equation (4.1) is not satisfied. We can define a higher product

$$\mu_d: \text{CM}^*(f_{d-1}, f_d) \otimes \dots \otimes \text{CM}^*(f_0, f_1) \rightarrow \text{CM}^*(f_0, f_d)$$

by linearly extending the formula

$$(4.2) \quad \Omega_{p_{d-1,d}} \otimes \dots \otimes \Omega_{p_{0,1}} \mapsto \sum_{p_{0,d}} n(p_{0,d}, p_{d-1,d}, \dots, p_{0,1}) \Omega_{p_{0,d}},$$

where each $\Omega_{p_{i,j}}$ is a choice of orientation of the appropriate ascending manifold. As in the case of the differential, it is clear from the construction of Appendix C that even though we have chosen orientations for the ascending manifolds in order to write a formula, the higher product is independent of all such choices. Modulo signs, Lemma 4.20 immediately implies the next result.

Proposition 4.22. *If \vec{f} is a set of functions which is Morse-Smale and boundary convex, then the A_∞ equation (3.3) holds.* \square

The proof of this proposition with signs is given in Appendix C.

We can now see that Morse functions on a manifold with boundary fit within Kontsevich and Soibelman’s framework of pre-categories.

Definition 4.23. *The **Morse pre-category** of Q is the pre-category $\text{Morse}(Q)$ defined by the data:*

- *Objects are functions on Q .*

- *Transverse sequences are sequences of functions which are Morse-Smale and boundary convex.*
- *Morphisms between a transverse pair (f, g) are given by the Morse complex $\text{CM}^*(f, g)$.*
- *Higher products are given by Equation (4.2).*

4.2. From Morse chains to cellular chains. The proof that ordinary and Morse homology agree extends to our setting.

Lemma 4.24.

$$\text{HM}^*(Q, f) \cong H^*(Q, \partial_f^+ Q)$$

Proof. Consider the union of all descending manifolds

$$(4.3) \quad W^s(f) = \bigcup_p \overline{W}^s(p).$$

Since f has no critical points outside $W^s(f)$, the image of any point in $Q - W^s(f)$ under the gradient flow intersects $\partial_f^+ Q$. In particular, we have a retraction

$$(4.4) \quad Q - W^s(f) \rightarrow \partial_f^+ Q \cap (Q - W^s(f)).$$

We conclude that the relative cohomology group

$$H^*(Q - W^s(f), \partial_f^+ Q \cap (Q - W^s(f)))$$

vanishes, and hence that the inclusion of $W^s(f)$ induces an isomorphism

$$(4.5) \quad H^*(Q, \partial_f^+ Q) \cong H^*(W^s(f), \partial_f^+ Q \cap W^s(f)).$$

It suffices therefore to prove that there is an isomorphism

$$(4.6) \quad H^*(W^s(f), \partial_f^+ Q \cap W^s(f)) \cong \text{HM}^*(Q, f).$$

To prove this, we note that the Morse-Smale condition guarantees that the closure of $\overline{W}^s(p)$ intersects $\overline{W}^s(q)$ only if the degree of q is smaller than that of p . In particular, we consider the filtration of $W^s(f)$ by the subsets $W^{s,k}(f)$ consisting of the closures of ascending manifolds of critical points whose degree is smaller than or equal to k . Every sub-quotient

$$\left(W^{s,k}(f) \cup \left(\partial_f^+ Q \cap W^s(f) \right) \right) / \left(W^{s,k-1}(f) \cup \left(\partial_f^+ Q \cap W^s(f) \right) \right)$$

is homeomorphic to a wedge of spheres indexed by the critical points of degree k . The result follows by induction on k , and the inductive step is a classical computation first appearing as Corollary 7.3 in Milnor's book [34] that, given a cobordism with critical points only in degree k and $k + 1$, the matrix of intersection pairings between ascending and descending manifolds defines a differential on this subquotient of the Morse complex whose cohomology agrees with the cohomology of the cobordism relative a certain subset of its boundary. \square

We would like to construct an explicit chain level map inducing the above isomorphism on cohomology. Further, we would like to relate the cup products at the chain level. This essentially entails the construction of A_∞ functors relating the Morse pre-category to a pre-category whose morphisms are chain complexes that compute ordinary cohomology. There are many possible options for this putative pre-category, such as singular cocycles, pseudo-cycles, or currents. If we're willing to work over \mathbb{R} , then currents would be the most convenient category to work with.

However, we would like to keep working over \mathbb{Z} while we still can. Unfortunately, pseudo-cycles offer the disadvantage that the intersection product is only partially defined (although, see [31] for an algebraic solution to this problem), while it is not clear to the author how to relate the tree product on the Morse complex to the products on singular cochains without using the abstract techniques of acyclic models as used by Guggenheim to relate the singular cup product with the wedge product on forms [22].

We will choose instead to work with the simplicial cochains of a fixed triangulation, although we will use cellular chains of dual triangulations as an intermediary. We will need to introduce a new moduli space of trees.

Definition 4.25. *The moduli space of shrubs \mathcal{S}_d is the moduli space of ribbon trees with d incoming edges all of which have finite type, whose outgoing edge has infinite type, and such that the distances from the node (interior vertex of the outgoing edge) to all the incoming vertices are equal.*

In Appendix B we construct a compactification $\overline{\mathcal{S}}_d$ of \mathcal{S}_d and prove the necessary smoothness and compactness results.

As with the ribbon trees which form the Stasheff polyhedra, we can endow the edges of a tree in \mathcal{S}_d with a natural orientation. Further, if \vec{f} is a sequence of $d+1$ functions, then every edge of a tree $T \in \mathcal{S}_d$ can be labeled uniquely by a difference $f_i - f_j$ satisfying the conditions of Definition 4.19.

Definition 4.26. *Let \vec{f} be a sequence of $d+1$ functions on Q , \vec{C} a collection of d cycles (C_1, \dots, C_d) , and p a critical point of $f_d - f_0$. The moduli space of **gradient trees from \vec{C} to p** , denoted $\mathcal{S}(p, \vec{C})$, consists of all maps*

$$\phi : T \rightarrow Q$$

with $T \in \mathcal{S}_d$ satisfying the following conditions:

- The image of the i th incoming vertex lies in C_i , and
- the outgoing vertex is mapped to p , and
- each edge e is a gradient line of the corresponding function f_e as in Definition 4.19.

Remark 4.27. The notation $\mathcal{S}(p, \vec{C})$ is ambiguous since the function \vec{f} whose gradient trees we are considering are elided. In practice, however, this should cause no confusion.

In Appendix B, we construct a decomposition of \mathcal{S}_d into smooth manifolds with corners corresponding to the different topological types of a shrub, yielding a decomposition of the closure $\overline{\mathcal{S}}_d$ into compact manifolds with corners. To construct an analogous decomposition of $\mathcal{S}(p, \vec{C})$, we consider evaluation maps

$$\text{ev}_i : W^s(p) \times \mathcal{S}_d \rightarrow Q$$

which are defined for each $1 \leq i \leq d$. Given a point (x, T) , its image under ev_i is simply the image of x under the composition of gradient flows corresponding to the path in T from the node of T to the input labeled by i . It is easy to see that $\mathcal{S}(p, \vec{C})$ is simply the inverse image of $C^1 \times \dots \times C^d$ under the map

$$\text{ev} : W^s(p) \times \mathcal{S}_d \rightarrow Q^d$$

whose i^{th} component is ev_i .

We will now fix a triangulation \mathcal{P} of Q such that ∂Q is a subcomplex. In general, since Morse-Smale functions on a disc can be arbitrarily complicated, there is no hope to recover any information about the Morse homology of an arbitrary function by relating it to the cellular complex of \mathcal{P} . We therefore fix a subdivision $\check{\mathcal{P}}$ dual to \mathcal{P} , and we shall include a condition on the behaviour of the Morse flow near strata of $\check{\mathcal{P}}$ as a restriction on the Morse functions we shall study.

In addition, we fix a family $\{\kappa_t\}_{t \in [0,1]}$ of diffeomorphisms of Q such that κ_0 is the identity, and such that $\check{\mathcal{P}}_t$, the subdivision obtained by applying κ_t to $\check{\mathcal{P}}$, is in simplicial position with respect to $\check{\mathcal{P}}_{t'}$ whenever $t < t'$ (see Appendix E). In particular, whenever t_i is an increasing sequence, the subdivisions $\check{\mathcal{P}}_{t_i}$ are mutually transverse.

Given a cell C or a sequence of cells $\vec{C} = (C_1, \dots, C_d)$ of $\check{\mathcal{P}}$, we will write C^t for the corresponding cell of $\check{\mathcal{P}}_t$ or $\vec{C}(\vec{t})$ for the sequence

$$(C_1^{t_1}, \dots, C_d^{t_d})$$

whenever $\vec{t} = (t_1, \dots, t_d)$ satisfies $0 \leq t_1 < \dots < t_d$.

We will further write, as in Appendix E, $\check{\mathcal{P}}_{t_1} \pitchfork \check{\mathcal{P}}_{t_2}$ for the cellular subdivision of Q consisting of the intersection of all cells in $\check{\mathcal{P}}_{t_1}$ and $\check{\mathcal{P}}_{t_2}$, and

$$\text{coll}_{t_2}: \check{\mathcal{P}}_{t_1} \pitchfork \check{\mathcal{P}}_{t_2} \rightarrow \check{\mathcal{P}}_{t_2}$$

for the corresponding cellular map.

Given two cells $C_1^{t_1}$ and $C_2^{t_2}$ we can use the family κ_t to construct a map

$$(4.7) \quad (C_1^{t_1} \cap C_2^{t_2}) \times [t_1, t_2] \rightarrow Q$$

whose restriction to $s_1 \in [t_1, t_2]$ identifies $C_1^{t_1} \cap C_2^{t_2}$ with $C_1^{s_1} \cap C_2^{t_2}$. This map realizes a homotopy between the identity and the collapse of $C_1^{t_1} \cap C_2^{t_2}$ onto its image under coll_{t_2} , producing a cell of one higher dimension which we write $K(C_1, C_2, t_1, t_2)$. More generally, if $\vec{t} = (t_1, \dots, t_d)$ is a sequence of increasing positive real numbers, let $K(\vec{t})$ denote the subset of points (s_1, \dots, s_{d-1}) in $[t_1, t_d] \times [t_2, t_d] \times \dots \times [t_{d-1}, t_d]$ satisfying the conditions

$$s_1 \leq s_2 \leq \dots \leq s_{d-1}.$$

$K(\vec{t})$ is a convex polytope; in dimension 2, for example, this polytope is shown in Figure 5.

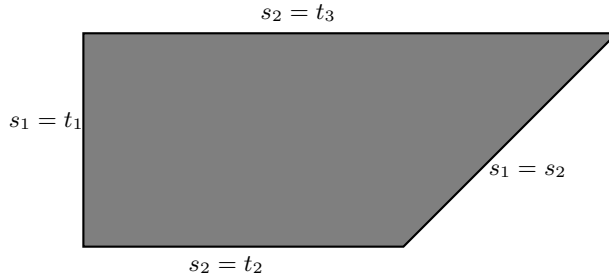


FIGURE 5.

Given a sequence of cells (C_1, \dots, C_d) of $\check{\mathcal{P}}$, Lemma E.5 implies that we have a diffeomorphism

$$(4.8) \quad C_1^{t_1} \cap \dots \cap C_{d-1}^{t_{d-1}} \cap C_d^{t_d} \cong C_1^{s_1} \cap \dots \cap C_{d-1}^{s_{d-1}} \cap C_d^{t_d}$$

that varies continuously with $\vec{s} \in K(\vec{t})$. In particular, we obtain a map

$$(4.9) \quad K(\vec{C}, \vec{t}) \cong \bigcap_i C_i^{t_i} \times K(\vec{t}) \rightarrow Q.$$

We now construct a new pre-category structure on $\text{Morse}(Q)$:

Definition 4.28. *A pair of functions (f_0, f_1) is transverse in $\text{Morse}(\check{\mathcal{P}})$ if and only if*

- *the pair (f_0, f_1) is transverse in $\text{Morse}(Q)$, and*
- *the inclusion*

$$\partial_f^+ Q \subset \bigcup_{\substack{C \in \check{\mathcal{P}} \\ C \cap \partial_f^+ Q \neq \emptyset}} C \equiv \partial_f^+ \check{\mathcal{P}}$$

is a deformation retract.

A sequence of functions (f_0, \dots, f_d) is transverse in $\text{Morse}(\check{\mathcal{P}})$ if and only if

- *it is transverse in $\text{Morse}(Q)$, and*
- *every subsequence consisting of two elements is transverse.*

The second condition required for the transversality of a pair of functions essentially forces the dynamics of the gradient flow of $f_1 - f_0$ to be no more (and no less) complex than the inclusion of $\partial_f^+ \check{\mathcal{P}}$ into $\check{\mathcal{P}}$.

In Appendix E, we construct a pre-category $\text{Cell}(\check{\mathcal{P}})$ with objects cell subdivisions which are C^1 -close to $\check{\mathcal{P}}$ (together with extra data at the boundary corresponding to the subsets $\partial_f^+ \check{\mathcal{P}}$), and where the cup product of simplicial cochains is realized as a geometric intersection pairing. We write $\text{Cell}_\kappa(\check{\mathcal{P}})$ for the subcategory of $\text{Cell}(\check{\mathcal{P}})$ consisting of cell subdivisions obtained by moving $\check{\mathcal{P}}$ along κ_t for some t . We now consider a new pre-category $\text{Cell} - \text{Morse}(\check{\mathcal{P}})$ which will serve as an intermediary between $\text{Cell}_\kappa(\check{\mathcal{P}})$ and $\text{Morse}(\check{\mathcal{P}})$. Again, we do this in two steps.

Definition 4.29. *The objects of $\text{Cell} - \text{Morse}(\check{\mathcal{P}})$ are pairs $(\check{\mathcal{P}}_{t_0}, f_0)$, with f_0 a function on Q .*

Two objects $(\check{\mathcal{P}}_{t_0}, f_0)$ and $(\check{\mathcal{P}}_{t_1}, f_1)$ are transverse if and only if

- (i) *t_1 is greater than t_0 , and*
- (ii) *(f_0, f_1) are transverse in $\text{Morse}(\check{\mathcal{P}})$, and*
- (iii) *the moduli spaces of shrubs with endpoints on a cell of $\check{\mathcal{P}}_{t_1}$ are regular.*

*If $(\check{\mathcal{P}}_0, f_0)$ and $(\check{\mathcal{P}}_1, f_1)$ are transverse objects of $\text{Cell} - \text{Morse}(\check{\mathcal{P}})$, the **space of morphisms** between them is defined to be*

$$C_{n-*}(\check{\mathcal{P}}_1 - \partial_{f_1 - f_0}^+ \check{\mathcal{P}}_1).$$

Since our ultimate goal is to pass from $\text{Cell} - \text{Morse}(\check{\mathcal{P}})$ to $\text{Morse}(\check{\mathcal{P}})$, we shall require transversality between the gradient flow lines of functions f_i and the cells of various subdivisions $\check{\mathcal{P}}_{t_i}$.

To state this formally, given a sequence of cells $\vec{C} = (C_1, \dots, C_d)$ of $\check{\mathcal{P}}$, times $\vec{t} = \{t_1 < \dots < t_d\}$, and a partition \vec{E} of the set $\{1, \dots, d\}$ into subset

$$(4.10) \quad \vec{E} = \{E_1, \dots, E_e\} = \{\{1, 2, \dots, i_2 - 1\}, \{i_2, \dots, i_3 - 1\}, \dots, \{i_e, \dots, d\}\},$$

we write $\vec{C}_{E_k} = (C_{i_k}, \dots, C_{i_{k+1}-1})$, and \vec{t}_{E_k} for the corresponding subsequence of \vec{t} . We consider

$$K_{\vec{E}}(\vec{C}, \vec{t}) = K(\vec{C}_{E_1}, \vec{t}_{E_1}) \times \dots \times K(\vec{C}_{E_e}, \vec{t}_{E_e}).$$

Generalising Definition 4.26, we have a moduli space

$$(4.11) \quad \mathcal{S}_e(p, K_{\vec{E}}(\vec{C}, \vec{t}))$$

obtained as a fiber product

$$(4.12) \quad \mathcal{S}_e(p, Q, \dots, Q) \times_{Q^e} K_{\vec{E}}(\vec{C}, \vec{t})$$

where the map from $K_{\vec{E}}(\vec{C}, \vec{t})$ to Q^e is given component-wise by Equation (4.9).

We incorporate a transversality conditions for these moduli spaces of shrubs with endpoints on $K_{\vec{E}}(\vec{C}, \vec{t})$ into our definition of the category $\text{Cell} - \text{Morse}(\check{\mathcal{P}})$. We shall need this in order to define a functor to $\text{Morse}(\check{\mathcal{P}})$.

Definition 4.30. A sequence $\{(\check{\mathcal{P}}_{t_i}, f_i)\}_{i=0}^d$ is transverse in $\text{Cell} - \text{Morse}(\check{\mathcal{P}})$ if

- (i) the sequence $\{\check{\mathcal{P}}_{t_i}\}_{i=0}^d$ is transverse in $\text{Cell}(\check{\mathcal{P}})$, and
- (ii) the sequence $\{f_i\}_{i=0}^d$ is transverse in $\text{Morse}(\check{\mathcal{P}})$, and
- (iii) if $\vec{C} = (C_1, \dots, C_d)$ is any sequence of cells in $\check{\mathcal{P}}$, \vec{t}' is a subsequence of \vec{t} of length d' with corresponding subsequence \vec{C}' of \vec{C} , and \vec{E}' a partition of $\{1, \dots, d'\}$ as in Equation (4.10), then the moduli spaces

$$\mathcal{S}_{e'}(p, K_{\vec{E}'}(\vec{C}', \vec{t}'))$$

are regular for any critical point p of $f_d - f_0$.

Note that whenever $\vec{C}' = \vec{C}$, and \vec{E} is the trivial partition $\{\{1\}, \{2\}, \dots, \{d\}\}$, the last conditions is imposing the regularity of the moduli spaces

$$\mathcal{S}_d(\vec{C}, \vec{t}).$$

Let us now consider the situation where a moduli space

$$\mathcal{S}_e(p, K_{\vec{E}}(\vec{C}, \vec{t}))$$

has virtual dimension 0. If the partition \vec{E} contains a subset E_k which consists of more than one element, then the evaluation map

$$K_{\vec{E}}(\vec{C}, \vec{t}) \rightarrow Q^d$$

has image which is contained in a neighbourhood of a cell of \mathcal{P}^d whose dimension is strictly smaller than that of $K_{\vec{E}}(\vec{C}, \vec{t})$. In particular, if all parameters t_i are sufficiently small, we expect this moduli space to be empty. Indeed, assuming regularity, the moduli problem for the cell of lower dimension has no solutions, so there should be no solution for C^0 close cells, regardless of their dimension. In this case, we could avoid introducing the higher dimensional cells $K_{\vec{E}}(\vec{C}, \vec{t})$ and define a more straightforward transversality condition using only the moduli spaces $\mathcal{S}_d(p, \vec{C}(\vec{t}))$ which would be essentially independent of sufficiently small choices of (t_1, \dots, t_d) . We elaborate further on this point in the next section when it is time to prove the A_∞ equation for the functor from the cellular to the Morse category.

The problem is that while the moduli spaces of shrubs with inputs on cells of \mathcal{P} are regular for generic choices of Morse functions, their compactifications do not satisfy this property because cells of \mathcal{P} do not necessarily intersect transversally. It seems likely that the moduli spaces $\mathcal{S}_e(p, K_{\vec{E}}(\vec{C}, \vec{t}))$ are still empty whenever some $|E_k| > 0$ for appropriate choices of diffeomorphisms κ_t , but proving this would require more subtle transversality arguments.

Lemma 4.31. *The virtual dimension of $\mathcal{S}_e(p, K_{\vec{E}}(\vec{C}, \vec{t}))$ is*

$$(d-1) + \deg(p) - \sum_i \operatorname{codim}(C_i).$$

In particular, it is independent of the subdivision \vec{E} . \square

4.2.1. *The functor defined by shrubs.* In this section, we will explain the proof of:

Proposition 4.32. *The pre-categories $\operatorname{Cell}(\check{\mathcal{P}})$ and $\operatorname{Morse}(\check{\mathcal{P}})$ are equivalent as A_∞ pre-categories.*

Before proceeding with the construction, let us make the auxiliary choice of an outward pointing vector field ν on ∂Q which is parallel to the normal vector of the boundary. We let

$$H(f) = \langle \operatorname{grad}(f), \nu \rangle$$

and observe that $H(f)$ is positive precisely on $\partial_f^+ Q$. The function $H(f)$ will play a rôle analogous to that of the function H in Appendix E. In particular, we have

$$\partial_{H(f_1-f_2)}^+ \check{\mathcal{P}}_1 = \partial_{f_1-f_2}^+ \check{\mathcal{P}}_1,$$

giving us two (slightly) incompatible notations.

We are now ready to use $\operatorname{Cell} - \operatorname{Morse}(\check{\mathcal{P}})$ for its stated purpose. First, observe that the map on objects

$$(\check{\mathcal{P}}_{t_0}, f_0) \rightarrow (\check{\mathcal{P}}_{t_0}, H(f_0))$$

induces the obvious equivalence of A_∞ -pre-categories from $\operatorname{Cell} - \operatorname{Morse}(\check{\mathcal{P}})$ to $\operatorname{Cell}_\kappa(\check{\mathcal{P}})$; the subcategory of $\operatorname{Cell}(\check{\mathcal{P}})$ consisting of subdivisions obtained by applying κ_t to \mathcal{P} for some κ_t . The inclusion

$$\operatorname{Cell}_\kappa(\check{\mathcal{P}}) \rightarrow \operatorname{Cell}(\check{\mathcal{P}})$$

is clearly essentially surjective.

We now construct an A_∞ functor from $\operatorname{Cell} - \operatorname{Morse}(\check{\mathcal{P}})$ to $\operatorname{Morse}(\check{\mathcal{P}})$. The functor is the forgetful map on objects

$$(\check{\mathcal{P}}_{t_0}, f_0) \rightarrow f_0.$$

The linear term of the functor is

$$\begin{aligned} \mathcal{F}_{cm} : C_{n-k}(\check{\mathcal{P}}_{t_1} - \partial_{f_1-f_0}^+ \check{\mathcal{P}}_{t_1}) &\rightarrow \operatorname{CM}^k(f_0, f_1) \\ C^{t_1} &\mapsto \sum_{p \in \operatorname{Crit}(f_1-f_0)} m(p, C^{t_1}) \Omega_p, \end{aligned}$$

where $m(p, C^{t_1})$ is the signed count of elements of $\mathcal{S}(p, C^{t_1})$. Note that the transversality assumptions in $\operatorname{Cell} - \operatorname{Morse}(\check{\mathcal{P}})$ ensure that $\mathcal{S}(p, C^{t_1})$ consists only of points, given that C^{t_1} has codimension k and p has degree k . The sign of a gradient trajectory is simply the sign of the intersection of the oriented chain C^{t_1} with the ascending submanifold $W^s(p)$ carrying the orientation Ω_p . To prove that \mathcal{F}_{cm} is a chain map, we need to understand 1-dimensional moduli spaces:

Claim. The union of the manifolds $\mathcal{S}(q, C^{t_1})$ over all critical points q of degree $k+1$ is a 1-dimensional manifold admitting a compactification whose boundary points correspond to rigid gradient trajectories from ∂C^{t_1} to q , or to broken gradient trajectories from C^{t_1} to q passing through a critical point p of degree k .

In the absence of a boundary on Q , this is a standard result about moduli spaces of gradient flow lines. The only difference for $\partial Q \neq \emptyset$ is that we must prove compactness of moduli spaces separately. However, transversality in $\text{Morse}(\check{\mathcal{P}})$ subsumes the boundary convexity of the function $f_1 - f_0$. In particular, gradient lines cannot escape to the boundary so the desired result holds. By interpreting the boundary terms algebraically, we conclude that

$$\mathcal{F}_{cm} \circ \partial = \mu_1 \circ \mathcal{F}_{cm}.$$

To see that \mathcal{F}_{cm} is a quasi-isomorphism, we recall that $m(p, C)$ is the signed intersection number between the oriented chain C^{t_1} and the ascending submanifold $W^s(p)$. Consider the complex of chains $C_*^{\text{tr}}(Q, \partial_{f_1-f_0}^+ Q)$ which are transverse to cells of $\check{\mathcal{P}}_{t_1}$. By assigning to each critical point its stable manifold, we obtain a map

$$\text{CM}_*(f_1 - f_0) \rightarrow C_*^{\text{tr}}(Q, \partial_{f_1-f_0}^+ Q),$$

which is a quasi-isomorphism by Lemma 4.24. By construction, \mathcal{F}_{cm} intertwines the duality pairing between Morse homology and cohomology, with the intersection pairing between $C_*^{\text{tr}}(Q, \partial_{f_1-f_0}^+ Q)$ and $C_{n-*}(\check{\mathcal{P}}_{t_1} - \partial_{f_1-f_0}^+ \check{\mathcal{P}}_{t_1})$. Since both pairings are known to be perfect pairings on homology, it follows that \mathcal{F}_{cm} must also be a quasi-isomorphism.

Let us describe the polynomial extensions of the functor \mathcal{F}_{cm} . Given a collection $\{(\check{\mathcal{P}}_{t_i}, f_i)\}_{i=0}^d$ of $d+1$ transverse objects in $\text{Cell} - \text{Morse}(\check{\mathcal{P}})$, and cycles $\{C_i^{t_i}\}_{i=1}^d = \vec{C}(\vec{t})$ belonging to $C_*(\check{\mathcal{P}}_{t_i} - \partial_{f_{i+1}-f_i}^+ \check{\mathcal{P}}_{t_i})$, and a partition E as in Equation (4.10), we consider a map

$$\mathcal{H}_E : C_*(\check{\mathcal{P}}_{t_d} - \partial_{f_{d+1}-f_d}^+ \check{\mathcal{P}}_{t_d}) \otimes \cdots \otimes C_*(\check{\mathcal{P}}_{t_1} - \partial_{f_1-f_0}^+ \check{\mathcal{P}}_{t_1}) \rightarrow \text{CM}^*(f_0, f_d)$$

whose action on generators is given by

$$C_d^{t_d} \otimes \cdots \otimes C_1^{t_1} \mapsto \sum_{p \in \text{Crit}(f_d - f_0)} m(p, K_{\vec{E}}(\vec{C}, \vec{t})) \Omega_p,$$

where $m(p, K_{\vec{E}}(\vec{C}, \vec{t}))$ is a signed count of the points in the moduli space

$$\mathcal{S}_e(p, K_{\vec{E}}(\vec{C}, \vec{t})).$$

We define

$$(4.13) \quad \mathcal{F}_{cm}^d = \sum_E \mathcal{H}_E.$$

An easy application of Lemma 4.31 implies that \mathcal{F}_{cm}^d has degree $1 - d$ as expected.

We must prove that the A_∞ equation for functors

$$(4.14) \quad \sum_{l_1 + \dots + l_k = d} \mu_k \circ (\mathcal{F}_{cm}^{l_1} \otimes \cdots \otimes \mathcal{F}_{cm}^{l_k}) = \sum_{i \leq d-1} (-1)^{\mathfrak{X}_i} \mathcal{F}_{cm}^d \circ (\mathbb{1}^{d-i-1} \otimes \partial \otimes \mathbb{1}^i) + \sum_{i \leq d-2} (-1)^{\mathfrak{X}_i} \mathcal{F}_{cm}^{d-1} \circ (\mathbb{1}^{d-i-2} \otimes \cup \otimes \mathbb{1}^i)$$

holds, where \mathfrak{X}_i is as in Lemma 3.3 and \cup stands for the cup product. In order to do this, we study, as usual, 1-dimensional moduli spaces of trees.

Before explaining the general case, we discuss why the first equation holds, ignoring all signs. This will shed light on the point of view that the term \mathcal{H}_E for

which E consists only of singletons is the essential part of the functor, and that the others terms are corrections imposed by the necessity to introduce perturbations of the subdivision $\check{\mathcal{P}}$. So we need to prove that:

(4.15)

$$\begin{aligned} \mu_2 \circ (\mathcal{H}_{\{1\}} \otimes \mathcal{H}_{\{2\}}) + \mu_1 \circ \mathcal{H}_{\{1\},\{2\}} + \mathcal{H}_{\{1\},\{2\}} \circ (\mathbb{1} \otimes \partial) + \mathcal{H}_{\{1\},\{2\}} \circ (\partial \otimes \mathbb{1}) \\ + \mathcal{H}_{\{2\}} \circ \cup = \mu_1 \circ \mathcal{H}_{\{1,2\}} + \mathcal{H}_{\{1,2\}} \circ (\mathbb{1} \otimes \partial) + \mathcal{H}_{\{1,2\}} \circ (\partial \otimes \mathbb{1}) \end{aligned}$$

Given two cells C_1 and C_2 such that the moduli space

$$\mathcal{S}_2(p, K_{\{1\},\{2\}}(\vec{C}, \vec{t})) = \mathcal{S}_2(p, C_1^{t_1}, C_2^{t_2})$$

is 1-dimensional, we recall that the boundary strata of the compactification correspond to the following possibilities

- (1) The distance between some internal point of the tree and an incoming vertex goes to infinity. Because of the restriction that the distance function to the incoming vertex is independent of the choice of incoming vertex, this can only happen if either (i) the outgoing edge converges to a broken gradient trajectory, or (ii) both incoming edges break. These two possibilities correspond to the terms:

$$\mu_2 \circ (\mathcal{H}_{\{1\}} \otimes \mathcal{H}_{\{2\}}) + \mu_1 \circ \mathcal{H}_{\{1\},\{2\}}$$

- (2) One of the incoming vertices escapes to the boundary of the cell that it lies on. This accounts for the terms

$$\mathcal{H}_{\{1\},\{2\}} \circ (\mathbb{1} \otimes \partial) + \mathcal{H}_{\{1\},\{2\}} \circ (\partial \otimes \mathbb{1}).$$

- (3) The length of an incoming edge can approach zero. By definition, both incoming edges collapse, so we end up with a gradient trajectory from $C_1^{t_1} \cap C_2^{t_2}$ to p .

This last type of boundary stratum is related, but not in general equal to, the term

$$\mathcal{H}_{\{2\}} \circ \cup.$$

Indeed, in Appendix E, we interpret the cup product as the composition of the intersection pairing with the collapse map

$$\text{coll}_{\mathcal{P}_{t_2}} : \check{\mathcal{P}}_{t_1} \pitchfork \check{\mathcal{P}}_{t_2} \rightarrow \check{\mathcal{P}}_{t_2}.$$

Note that the image of $C_1^{t_1} \cap C_2^{t_2}$ under the collapse map is $C_1^{t_2} \cap C_2^{t_2}$. We must therefore account for the difference between the number of gradient trajectories starting at $C_1^{t_1} \cap C_2^{t_2}$ and $C_1^{t_2} \cap C_2^{t_2}$, which we claim is precisely given by the remaining terms

$$\mu_1 \circ \mathcal{H}_{\{1,2\}} + \mathcal{H}_{\{1,2\}} \circ (\mathbb{1} \otimes \partial) + \mathcal{H}_{\{1,2\}} \circ (\partial \otimes \mathbb{1}).$$

Consider the 1-dimensional moduli space

$$\mathcal{S}_2(p, K_{\{1,2\}}(\vec{C}, \vec{t})) = \mathcal{S}_1(p, K(C_1, C_2, t_1, t_2))$$

of gradient trees with inputs lying on the image of the evaluation map from $K(C_1, C_2, t_1, t_2)$ and output the critical point p . By considering the boundary of this moduli space, we will find five terms respectively corresponding to the composition $\mathcal{H}_{\{2\}} \circ \cup$, the count of gradient trajectories starting at $C_1^{t_1} \cap C_2^{t_2}$, and the right hand side of Equation (4.15).

The compactification has two types of boundaries:

- (1) either the gradient trajectory breaks at a critical point q , corresponding to the term

$$\mu_1 \circ \mathcal{H}_{\{1,2\}},$$

- (2) or the gradient trajectory escapes to the boundary of $K(C_1, C_2, t_1, t_2)$.

By construction, the boundary of $K(C_1, C_2, t_1, t_2)$ can be decomposed geometrically as a union

$$(C_1^{t_1} \cap C_2^{t_2}) \bigcup K(\partial C_1, C_2, t_1, t_2) \bigcup K(C_1, \partial C_2, t_1, t_2) \bigcup (C_1^{t_2} \cap C_2^{t_1}).$$

The count of gradient trajectories starting on these 4 different strata completes the proof of the A_∞ equation for $d = 2$.

We now explain the general idea behind the proof of the A_∞ relation for functors which appears in Equation (4.14). We again consider \vec{C} and p such that

$$\mathcal{S}_e(p, K_{\vec{E}}(\vec{C}, \vec{t}))$$

has dimension 1. Let us recall the reasons why this moduli space is not compact:

- (1) The length of some internal edges can go to infinity. In 1-parameter families, the only possibility is that every path from an incoming vertex to the outgoing vertex contains precisely one edge whose length becomes infinite. This corresponds to the left hand side of Equation (4.14):

$$(4.16) \quad \sum_{\substack{k \\ l_1 + \dots + l_k = d}} \mu_k \circ (\mathcal{F}_{cm}^{l_1} \otimes \dots \otimes \mathcal{F}_{cm}^{l_k})$$

- (2) The length of an incoming edge can approach zero. As we are assuming transversality, this occurs for a pair of incoming edges (in particular, the collapse of triples of incoming edges does not occur in one-parameter families, and hence does not affect the differential). This corresponds to the count of elements in the moduli space

$$(4.17) \quad \mathcal{S}_{e-1} \left(p, K_{E_1}(\vec{C}_{E_1}, \vec{t}_{E_1}), \dots, K_{E_{i-1}}(\vec{C}_{E_{i-1}}, \vec{t}_{E_{i-1}}) \cap K_{E_i}(\vec{C}_{E_i}, \vec{t}_{E_i}), \dots, K_{E_d}(\vec{C}_{E_d}, \vec{t}_{E_d}) \right)$$

for some fixed i .

- (3) One of the incoming vertices escapes to the boundary of the cell $K_{E_j}(\vec{C}_{E_j}, \vec{t}_{E_j})$ that it lies on. This cell has three types of boundary strata. To simplify our notation, we consider the case where the partition consists of only one set E , allowing us to drop the subscript j from the discussion below:

- (a) We take the boundary of a cell in \vec{C}_E . This accounts for the first term in the right hand side of Equation (4.14):

$$(4.18) \quad \sum_{i \leq d-1} (-1)^{\mathfrak{A}_i} \mathcal{F}_{cm}^d \circ (\mathbb{1}^{d-i-1} \otimes \partial \otimes \mathbb{1}^i)$$

- (b) we can reach the boundary of a face in $K(\vec{t}_E)$ where some $s_i = s_{i+1}$, with the convention that $s_d = t_d$. Note that the corresponding cell in Q is the same as the one obtained by applying K to the sequence of cells

$$(C_1^{t_1}, \dots, C_i^{t^{i+1}} \cap C_{i+1}^{t^{i+1}}, \dots, C_d^{t^d})$$

whenever $C_i^{t^i}$ and $C_{i+1}^{t^{i+1}}$ intersect non-trivially. In this case, Lemma E.6 implies that the intersection $C_i^{t^i} \cap C_{i+1}^{t^{i+1}}$ is the result of applying the cup product to $C_i^{t^i}$ and $C_{i+1}^{t^{i+1}}$, so this boundary stratum corresponds to the remaining term of Equation (4.14):

$$(4.19) \quad \sum_{i \leq d-2} (-1)^{\mathfrak{K}_i} \mathcal{F}_{cm}^{d-1} \circ (\mathbb{1}^{d-i-2} \otimes \cup \otimes \mathbb{1}^i)$$

- (c) the remaining boundary stratum correspond to $s_i = t_i$ for some i . One can check that these boundary strata exactly cancel out with the strata coming from Equation (4.17).

The A_∞ equation will then hold for the usual argument that the algebraic count of the number of boundary points of a 1-dimensional manifold vanishes. To complete the argument, one should check the correctness of the signs. We omit the long, but straightforward, computation which can be performed along the lines of the arguments given in Appendix C to prove the correctness of the signs in the A_∞ -structure of $\text{Morse}(\check{\mathcal{P}})$.

Combining this result with Proposition E.10, we conclude

Corollary 4.33. *The pre-categories $\text{Simp}(\mathcal{P})$ and $\text{Morse}(\check{\mathcal{P}})$ are equivalent as A_∞ pre-categories.* \square

4.3. The equivalence of Fukaya and Morse pre-categories. Given a collection of functions (f_0, \dots, f_n) on a smooth compact Riemannian manifold we obtain a collection of Lagrangians L_i , the graphs of df_i , in T^*M , and the critical points $x_i \in \text{Crit}(f_i - f_j)$ correspond to the intersection points of L_i with L_j . This yields a canonical isomorphism between the underlying graded vector spaces of Floer and Morse complexes. In [17], Fukaya and Oh extended Floer's results that the differentials on Floer and Morse theory agree. These results have been reworked and amplified in [9, 37]. Although the results are usually stated for manifolds without boundary, one can easily extend them to the case where the boundary is not empty. In our notation below, ϵg stands for the rescaled metric, and ϵJ for the induced almost complex structure on T^*Q .

Proposition 4.34. *Let (Q, g) be a compact Riemannian manifold possibly with boundary. For every generic finite collection of functions (f_0, \dots, f_n) which are transverse in $\text{Morse}(Q)$ and which have no boundary critical points, there exists a constant $\epsilon_0 > 0$ such that for all $\epsilon < \epsilon_0$ and for every collection*

$$\{x_{0,n}, x_{0,1}, \dots, x_{n-1,n}\}$$

of critical points $x_{i,j}$ of $f_j - f_i$, there exists an orientation preserving homeomorphism between moduli spaces of holomorphic discs and gradient trees:

$$\mathcal{M}_{\epsilon J}(x_{0,n}; x_{0,1}, \dots, x_{n-1,n}) \cong \mathcal{T}_{\epsilon g}(x_{0,n}; x_{0,1}, \dots, x_{n-1,n}).$$

Further, every holomorphic curves is C^0 close to the corresponding gradient tree

Proof. Extend Q to a smooth compact manifold without boundary \hat{Q} , and extend the function f_i to smooth functions on \hat{Q} . By Fukaya-Oh's result, there exists an ϵ , small enough, so that moduli spaces of holomorphic discs and gradient trees on \hat{Q} are in bijective correspondence, and all holomorphic discs are C^0 close to gradient trees. We write $\mathcal{N}_\delta(K)$ for the δ -neighbourhood of a set K in \hat{Q} . Note that the assumption that (f_0, \dots, f_n) are transverse in $\text{Morse}(Q)$ implies that, for

sufficiently small δ , all gradient trees contained in Q are contained in $Q - \mathcal{N}_\delta(\partial Q)$. Choosing ϵ small enough therefore guarantees that all holomorphic discs which project to $\mathcal{N}_\delta(Q)$ project to $Q - \mathcal{N}_\delta\partial Q$. We conclude that for ϵ small enough, the theorem of Fukaya-Oh holds for the manifold with boundary Q . \square

Remark 4.35. If the sequence \vec{f} is allowed to have critical points at the boundary, then boundary convexity implies that there is some neighbourhood of the boundary in which gradient trees of interior critical points do not enter. For sufficiently small ϵ , the argument above implies that count of holomorphic discs in the complement of this neighbourhood agrees with the count of gradient trees.

Fukaya and Oh also state

Corollary 4.36. *The Fukaya pre-category of exact Lagrangian sections of T^*Q is A_∞ quasi-isomorphic to the Morse pre-category of Q .*

We will write $\text{Fuk}(Q)$ for this category of exact Lagrangian sections. In order for this to make sense when Q is a manifold with boundary, we need to make sure that we have compactness for moduli spaces of holomorphic curves. In the application we have in mind, this follows from an extra structure imposed at infinity (admissibility with respect to the “superpotential”). We explain the (completely formal) proof of the above Corollary in the absence of a boundary.

Proof of Corollary 4.36. Let $\epsilon_i \rightarrow 0$ be a sequence of rescaling parameters, and let

$$\text{Fuk}(Q; \epsilon_i J)$$

be the corresponding Fukaya categories. The homotopy method (see [16, 40]) defines A_∞ equivalences of pre-categories

$$\Phi_{i,i+1}: \text{Fuk}(Q; \epsilon_i J) \rightarrow \text{Fuk}(Q; \epsilon_{i+1} J),$$

which is essentially the identity on objects. In fact, because of transversality problems, this functor is only defined on a full pre-subcategory of $\text{Fuk}(Q; \epsilon_i J)$ consisting of objects for which the family of almost complex structure $\epsilon_i J$ satisfies an appropriate transversality condition. For simplicity of notation, we will neglect this problem.

For each i , we consider the sub-pre-category $\text{Fuk}_i(Q; \epsilon_i J)$ where a collection (L_0, \dots, L_n) of Lagrangians is transverse if and only if the moduli spaces

$$\mathcal{M}_{\epsilon_i J}(x_{0,n}; x_{0,1}, \dots, x_{n-1,n}) \cong \mathcal{T}_{\epsilon_i g}(x_{0,n}; x_{0,1}, \dots, x_{n-1,n}).$$

are diffeomorphic for all intersection points $x_{i,j} \in L_i \cap L_j$ and for all $\epsilon < \epsilon_i$. The functors $\Phi_{i,i+1}$ preserve these subcategories in the sense that if a sequence is transverse in $\text{Fuk}_i(Q; J_{\epsilon_i})$, its image under $\Phi_{i,i+1}$ is transverse in $\text{Fuk}_{i+1}(Q; \epsilon_{i+1} J)$. In particular, we obtain a directed system (in fact, a nested embedding)

$$(4.20) \quad \text{Fuk}_0(Q; \epsilon_0 J) \rightarrow \Phi_{0,1}^{-1}(\text{Fuk}_1(Q; \epsilon_1 J)) \rightarrow \Phi_{0,2}^{-1}(\text{Fuk}_2(Q; \epsilon_2 J)) \rightarrow \dots$$

whose limit is $\text{Fuk}(Q; J_{\epsilon_0})$. We can think of this as a filtration of $\text{Fuk}(Q; J_{\epsilon_0})$.

Denote by j_i the “identity functor”

$$\text{Fuk}_i(Q; \epsilon_{i+1} J) \rightarrow \text{Morse}_g(Q)$$

taking every exact Lagrangian section to its defining function. By construction, there is an equality of (restricted) functors

$$j_0 = j_1 \circ \Phi_{0,1}$$

on $\text{Fuk}_0(Q; \epsilon_0 J)$. This implies that we can define the right hand side as the extension of j_0 to the subcategory $\Phi_{0,1}^{-1}(\text{Fuk}_1(Q; \epsilon_1 J))$. Proceeding inductively, we can extend j_0 through every step of the direct limit of Equation (4.20). Passing to the limit, we obtained the desired A_∞ equivalence

$$\text{Fuk}_0(Q; \epsilon_0 J) \rightarrow \text{Morse}_g(Q)$$

□

5. HMS FOR TORIC VARIETIES: FUKAYA = MORSE

Having established the general framework for relating the multiplicative structures on various chain models of ordinary homology, we return to our goal of proving our main result on homological mirror symmetry for toric varieties.

5.1. Tropical Lagrangians as a Morse pre-category. The observation of Section 3.2 that all tropical Lagrangian sections lift to graphs of exact 1-forms allows us to define a Morse pre-category.

Definition 5.1. *The Morse pre-category of \mathcal{Q} is the pre-category $\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$ given by the following data*

- *Objects are functions on \mathcal{Q} the graphs of whose differentials agree with lifts to $T^*\mathbb{R}^n$ of tropical Lagrangian sections in $((\mathbb{C}^*)^n, M)$.*
- *Transverse sequences are given by functions (f_0, f_1, \dots, f_d) which are Morse-Smale and define lifts of a positive sequence of Lagrangians (L_0, \dots, L_d) .*
- *Morphisms between f_0 and f_1 on \mathcal{Q} are given by*

$$\text{CM}_{\mathbb{Z}^n}^*(f_0, f_1) = \bigoplus_{u \in \mathbb{Z}^n} \text{CM}^*(f_0, f_1 + u).$$

- *Higher products are defined by Equation (4.2).*

Remarks 5.2. As in the previous section, the graded abelian group $\text{CM}^*(f_0, f_1 + u)$ is generated by interior critical points. Further, when we write $f_1 + u$, we're thinking of u as a linear function on \mathbb{R}^n

$$y \mapsto \langle y, u \rangle.$$

We have chosen this pre-category for its convenience. Its morphisms are \mathbb{Z}^n graded, and the higher products respect this grading. If we further required quasi-isomorphisms to live in the 0 grading, then we would obtain a pre-category which is equivalent to the category of equivariant $(\mathbb{C}^*)^n$ line bundles on the mirror toric variety. However, we will not impose such a restriction, so that different lifts of the same function lie in the same quasi-isomorphism class. We could have defined the objects to be tropical Lagrangian sections, but our spaces of morphisms would have only been graded by an affine space over \mathbb{Z}^n .

Lemma 5.3. *$\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$ is a unital A_∞ pre-category.*

Proof. We must prove that the A_∞ equation holds for transverse sequences. By the results of Section 4.1, this follows from the

Claim. The lifts of a sequence of admissible Lagrangians satisfy the boundary convexity property of Lemma 4.20.

Proof of Claim. It suffices to prove the results in a neighbourhood of each cell of $\partial\mathcal{Q}$. We first consider the cells labeled by vertices of ∂Q . On these open sets, the gradient vector fields of functions defining tropical Lagrangian sections are locally constant when restricted to the part of the boundary which lies in such neighbourhoods. Since \mathcal{Q} is a convex subset of \mathbb{R}^n , the result follows near the 0-skeleton. We now proceed inductively by using, near each i cell, the product decomposition which was defined in [1, Lemma 4.4]. \square

It remains to prove unitality, i.e. to produce a sufficient number of quasi-isomorphic copies of each object. We observed in [1, Lemma 5.8] that if $r_{\partial\mathcal{Q}}$ is any function which agrees with the square of the distance to the boundary in some neighbourhood of $\partial\mathcal{Q}$ then, for any real number K , the image of the graph of $-Kdr_{\partial\mathcal{Q}}$ under the map $W : (\mathbb{C}^*)^n \rightarrow \mathbb{C}$ is well approximated, near the origin, by a straight line of slope K . More generally, if f is a function whose differential defines an admissible tropical Lagrangian section, then for every real number K ,

$$f - Kr_{\partial\mathcal{Q}}$$

is C^2 close in some neighbourhood of $\partial\mathcal{Q}$ to a function f'_K defining another admissible tropical Lagrangian section such that the pair (f, f'_K) is positive (See [1, Lemma 4.15]). More precisely, if the curve corresponding to f has slope a at the origin, then the curve corresponding to f'_K has slope $a - K$,

Since

$$\mathrm{CM}^*(f, f'_K) \cong \mathrm{CM}^*(0, r) \cong H^*(\mathbb{R}^n),$$

we have a canonical element in $e \in \mathrm{CM}^0(f, f'_K)$ which represents the generator of $H^0(\mathbb{R}^n)$. By Proposition 4.32, this generator acts as a cohomological unit. In particular, if (f_0, \dots, f_n) is a finite transverse sequence of admissible functions, there exists a constant $K \gg 0$ such that the sequence

$$(f'_{-K}, f_0, \dots, f_n, f'_K)$$

defines a positive sequence of admissible Lagrangians. The constant K can simply be chosen to be much larger than the absolute value of the slopes at the origin of all the curves corresponding to the admissible Lagrangians defined by f_i and f . We can then choose a C^∞ -small perturbation of r away from the boundary of \mathcal{Q} to achieve transversality. \square

Remark 5.4. We were careful to state the results of Section 4.1, allowing critical points on the boundary in order to be able to apply them directly here. There is a slightly subtle point concerning the positivity condition. Consider the case where we are computing the space of morphisms between the zero-section and a Lagrangian defined by a function f . Whenever f has a boundary critical point, the set-up of Section 4.1 allows for this point to be included in $\partial_f^+ \mathcal{Q}$ or $\partial_f^- \mathcal{Q}$ depending on the behaviour of the gradient flow nearby. However, the positivity assumption precisely eliminates this ambiguity. Indeed, the gradient vector field of f points inward as soon as we move away from the boundary towards the interior. This implies that $\partial_f^- \mathcal{Q}$ is exactly the subset of $\partial\mathcal{Q}$ where the gradient flow of f satisfies the (non-strict) inequality

$$\langle \mathrm{grad}(f), \nu \rangle \leq 0.$$

5.2. Equivalence of Morse and Fukaya categories. In order to prove the equivalence of Morse and Fukaya categories, we need to establish more control over the tropical construction. While both Floer and Morse trajectories do not escape to the boundary, we cannot directly apply the Fukaya-Oh argument for two reasons:

- (1) The almost complex structure we're using near $\partial T^*\mathcal{Q}$ does not agree with the complex structure J induced by the flat metric on \mathcal{Q} (which is the restriction of the complex structure on $(\mathbb{C}^*)^n$), and
- (2) admissibility relies on the existence of a holomorphic function near the boundary. Such a function will not remain holomorphic upon “rescaling” J as in the Fukaya-Oh procedure.

We return to the notation where we use a semi-tropical degeneration producing a family of functions $W_{t,s}$, which define symplectic hypersurfaces $M_{t,s} \subset (\mathbb{C}^*)^n$, and hence a family of subsets $\mathcal{Q}_{t,s} \subset \mathbb{R}^n$ and almost complex structures $J_{t,s}$ on $(\mathbb{C}^*)^n$. We will also be using the notation of Section 4.3 so we write ϵJ for the almost complex structure induced on $(\mathbb{C}^*)^n = T^*\mathbb{R}^n/\mathbb{Z}^n$ by rescaling the metric on the base \mathbb{R}^n by ϵ . We will write \mathcal{L}^t for the part of the zero section of $(\mathbb{C}^*)^n$ lying over $\mathcal{Q}_{t,1}$, and $\tilde{\mathcal{L}}^t$ for its lift to the cotangent bundle.

In Appendix A we explain how to modify the construction given in [1] in order to prove the following result. We simply state and prove the case with one Lagrangian, the techniques we use readily extend to establish compactness theorems for finitely many Lagrangians.

Proposition 5.5. *Let L be an admissible tropical Lagrangian section with boundary on $M_{t_0,1}$ such that (\mathcal{L}^{t_0}, L) is a positive pair, and let f be a function defining a lift \tilde{L} of L to $T^*\mathbb{R}^n$. There exists a family f^t for $t \in [t_0, +\infty)$ of functions defining admissible tropical Lagrangian sections L^t with boundary on $M_{t,1}$ such that (\mathcal{L}^t, L^t) is positive, all interior intersection points between L^t and the zero section are transverse, and for sufficiently large t there are positive constants δ_1^t , δ_2^t and δ_3^t such that:*

- (1) *for any $u \in \mathbb{Z}^n$, the Morse gradient flows of $f - u$ connecting interior critical points lie in the complement of the δ_1^t neighbourhood of $\partial\mathcal{Q}_{t,1}$, and*
- (2) *all $J_{t,s}$ holomorphic discs with marked points on interior intersection points of L^t with the zero section project to the complement of the δ_2^t neighbourhood of $\partial\mathcal{Q}_{t,1}$, and*
- (3) *families of ϵJ holomorphic discs with marked points on interior intersection points of L^t and \mathcal{L}^t which for some ϵ_0 lie outside the $2\delta_3^t$ neighbourhood of $\partial\mathcal{Q}_{t,1}$ remain outside the δ_3^t neighbourhood of $\partial\mathcal{Q}_{t,1}$ for all ϵ , and*
- (4) *$2\delta_3^t$ is smaller than δ_1^t and δ_2^t .*

Assume that we have chosen a family L^t and t large enough to satisfy the above conditions. We will want to produce a chain isomorphism

$$\mathrm{CF}^*(\tilde{\mathcal{L}}^{t_0}, \tilde{L}; J_{t_0,1}) \rightarrow \mathrm{CM}^*(0, f).$$

While it is possible to do this in one step it will be conceptually easier for us to construct a sequence of chain isomorphisms

$$\begin{aligned} \mathrm{CF}^*(\tilde{\mathcal{L}}^{t_0}, \tilde{L}; J_{t_0,1}) &\rightarrow \mathrm{CF}^*(\tilde{\mathcal{L}}^t, \tilde{L}^t; J_{t,1}) \rightarrow \mathrm{CF}^*(\tilde{\mathcal{L}}^t, \tilde{L}^t; J_{t,0}) \\ &\rightarrow \mathrm{CM}^*(0, f_t) \rightarrow \mathrm{CM}^*(0, f). \end{aligned}$$

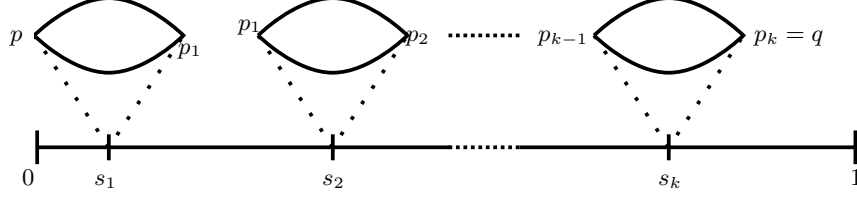


FIGURE 6.

All these maps can be constructed in essentially the same way via either continuation maps or the homotopy method, so we will focus on constructing a chain isomorphism

$$\mathcal{F}_{\mathbf{J}}: \text{CF}^*(\tilde{\mathcal{L}}^t, \tilde{L}^t; J_{t,1}) \rightarrow \text{CF}^*(\tilde{\mathcal{L}}^t, \tilde{L}^t; J_{t,0}).$$

Part (2) of Proposition 5.5 ensures that the right hand side is a complex; it states that holomorphic discs with respect to $J_{t,0}$ stay away from the boundary so we can apply the usual compactness arguments. We now introduce the relevant moduli spaces for the definition of $\mathcal{F}_{\mathbf{J}}$. The technique we are describing is originally due to Floer, although it has been extended by Fukaya, Oh, Ohta, and Ono in their work on Lagrangian Floer homology, [16]

If p and q are intersection points of \tilde{L}^t and $\tilde{\mathcal{L}}^t$, we write $\mathcal{M}_{\mathbf{J}}(p, q)$ for the union of all moduli spaces $\mathcal{M}_{J_{t,s}}(p, q)$ for fixed t and varying $s \in [0, 1]$. This space is naturally topologised as the vanishing locus of a section of an infinite dimensional vector bundle over the product of the interval $[0, 1]$ with the space of smooth maps from the disc to $T^*\mathbb{R}^n$ satisfying appropriate boundary conditions. As in the usual Floer theoretic setup, this allows us to define regularity as the requirement that the linearised operator is surjective at the zeros; i.e. that the section intersect the zero level set transversely. Of course, to make this discussion formal, one introduces appropriate Sobolev spaces. We refer the interested reader to the references [16, 40], and simply note that regardless of the details, it is easy to see that,

$$\text{virdim}(\mathcal{M}_{\mathbf{J}}(p, q)) = \text{virdim}(\mathcal{M}_{J_{t,s}}(p, q)) + 1.$$

We note that the condition that $\mathcal{M}_{\mathbf{J}}(p, q)$ is regular does not imply that each moduli space $\mathcal{M}_{J_{t,s}}(p, q)$ is regular. Indeed, we will consider moduli spaces $\mathcal{M}_{\mathbf{J}}(p, q)$ whose virtual dimension is 0. Assuming regularity, this space will therefore consist of finitely many J_{t,s_i} holomorphic discs which occur for some $s_i \in (0, 1)$. Since the virtual dimension of $\mathcal{M}_{J_{t,s}}(p, q)$ is -1 for any s , this moduli space would be empty if it were regular. We will often refer to these discs (which are only regular in families) as **exceptional discs**.

Using the usual transversality arguments, we see that for a generic choice of L^t :

- (1) all moduli spaces $\mathcal{M}_{J_{t,0}}(p, q)$ and $\mathcal{M}_{J_{t,1}}(p, q)$ are regular, and
- (2) the parametrised moduli space $\mathcal{M}_{\mathbf{J}}(p, q)$ is regular.

In this case, we define a new moduli space $\mathcal{H}(p, q)$ consisting of sequences of $J_{t,s}$ -holomorphic discs as in Figure 6. To be more precise:

Definition 5.6. *An element of $\mathcal{H}(p, q)$ consists of the following data:*

- (1) A collection $\{p = p_0, p_1, \dots, p_k = q\}$ of interior intersection points between \tilde{L}^t and the zero section, together with an increasing sequence $0 < s_1 < \dots < s_k < 1$, and

- (2) for each $0 < i \leq k$, a $J_{t,1-s_i}$ holomorphic disc u_i with boundary on $\tilde{L}^t \cup \tilde{\mathcal{L}}^t$ with the incoming point mapping to p_{i-1} and outgoing marked point mapping to p_i .

We note that the dimension of $\mathcal{H}(p, q)$ is given by the difference between the degrees of p and q . In particular, $\mathcal{H}(p, q)$ is discrete precisely whenever $\deg(p) = \deg(q)$ so we can define a degree 0 map

$$\begin{aligned} \mathcal{F}_{\mathbf{J}}: CF^*(\tilde{\mathcal{L}}^t, \tilde{L}^t; J_{t,1}) &\rightarrow CF^*(\tilde{\mathcal{L}}^t, \tilde{L}^t; J_{t,0}) \\ p &\mapsto p + \sum_{\deg(p)=\deg(q)} |\mathcal{H}(p, q)| q \end{aligned}$$

where $|\mathcal{H}(p, q)|$ is the signed count of the number of points in $\mathcal{H}(p, q)$. The following Lemma essentially goes back to Floer:

Lemma 5.7. *The map $\mathcal{F}_{\mathbf{J}}$ is a chain isomorphism.* \square

We briefly discuss the proof. First, we recall that the exactness condition implies that we have an energy filtration on Floer cochain groups: in our case the energy of an intersection point p is the value of f at p . By ordering all intersection points according to their action, we obtain a basis for the Floer cochain groups. The usual argument involving Stokes' theorem implies that the matrix representing $\mathcal{F}_{\mathbf{J}}$ in this basis is upper triangular with 1 along the diagonal. In particular $\mathcal{F}_{\mathbf{J}}$ will necessarily be a chain isomorphism if it happens to be a chain map.

To prove that $\mathcal{F}_{\mathbf{J}}$ is a chain map, one considers the compactification of the 1-dimensional moduli spaces $\mathcal{H}(p, q)$; this corresponds to the condition that $\deg(q) = \deg(p) + 1$. One can check that the only sequences $\{p = p_0, p_1, \dots, p_k = q\}$ which contribute to $\mathcal{H}(p, q)$ are those for which there is $i < k$ such that $\deg(p_j) = \deg(p)$ for all $j \leq i$ and all other intersection points have degree equal to $\deg(q)$. In this case, we have a collection of exceptional solutions together with one (regular) $J_{t,1-s_i}$ -holomorphic strip connecting p_i and p_{i+1} , and it is clear that we can allow s_i to vary, giving us precisely a 1-dimensional moduli space.

There is a natural bordification $\overline{\mathcal{H}}(p, q)$ of $\mathcal{H}(p, q)$ which allows broken curves, as well as the possibility that $s_0 = 0$ or $s_1 = 1$. The proof that $\mathcal{F}_{\mathbf{J}}$ is a chain map has now been reduced to an analysis of the boundary of the bordification $\overline{\mathcal{H}}(p, q)$. Two of the boundary strata correspond to having $s_0 = 0$, and $\deg(p_1) = \deg(p) + 1$ or, alternatively, $s_1 = 1$ and $\deg(p_{k-1}) + 1 = \deg(q)$. It is rather clear that these terms correspond to

$$\mu_1 \circ \mathcal{F}_{\mathbf{J}} + \mathcal{F}_{\mathbf{J}} \circ \mu_1.$$

We must therefore argue that $\overline{\mathcal{H}}(p, q)$ is compact, and that the contributions of all other boundary strata vanish. Using part (2) of Proposition 5.5, we see that $J_{t,s}$ holomorphic curves do not approach the boundary, so we can apply the usual compactness arguments to $\overline{\mathcal{H}}(p, q)$. As to the other boundary strata, we observe that they correspond to a broken curve consisting of two components, one exceptional and the other regular, occurring at some fixed time s_i . Note that if u_i is a regular curve at time s_i , and u_{i+1} is an exceptional curve at time s_{i+1} , we obtain exactly such a broken curve if we consider a family where s_i converges to s_{i+1} . However, the usual gluing procedure yields another family of curves converging to this broken curve. Taken together, these contributions cancel, proving that $\mathcal{F}_{\mathbf{J}}$ is a chain map.

The homotopy method was extended by Fukaya, Oh, Ohta, and Ono in their joint work [16] to produce A_∞ -functors; an account in the exact setting appears in [40, Section 10e]. The end result is that, using the analogue of Proposition 5.5 in the presence of more than one Lagrangian, we can conclude that the pre-categories

$$\mathcal{T}\text{Fuk}((\mathbb{C}^\star)^n, M_{t,1}; J_{t,1}), \mathcal{T}\text{Fuk}((\mathbb{C}^\star)^n, M_{t,1}; J_{t,0}).$$

are A_∞ equivalent.

Continuing along the same train of thought, we consider the family of almost complex structures ϵJ . Part (3) of Proposition 5.5 implies a compactness theorem for families of ϵJ -holomorphic curves which for any ϵ_0 project to the complement of the δ_3^t neighbourhood of $\partial\mathcal{Q}_{t,1}$. The assumption that $2\delta_3^t < \delta_2^t$ implies an equivalence between the chain-level multiplicative structures defined using (i) J -holomorphic curves, and (ii) those ϵJ holomorphic curves whose projecting to $\mathcal{Q}_{t,1}$ is contained in the complement of the δ_3^t neighbourhood of $\partial\mathcal{Q}_{t,1}$.

Further, by Fukaya and Oh's theorem, the moduli spaces of ϵJ holomorphic discs correspond, for ϵ small enough, to gradient trajectories of f^t . Since $\delta_3^t < \delta_1^t$, it is sufficient to count holomorphic curves which project to the complement of the δ_3^t neighbourhood of $\partial\mathcal{Q}_{t,1}$. Running time backwards (and applying the homotopy method one last time) we can relate the count of gradient trajectories for f^t to the one for the original function f .

To implement these ideas in the language of pre-categories, one has to introduce an intermediate pre-category whose objects are families of function f^t satisfying the conditions of Proposition 5.5, and whose morphisms are the Floer complexes at time $t = t_0$. By forgetting the family, we obtain an A_∞ quasi-equivalent to $\mathcal{T}\text{Fuk}((\mathbb{C}^\star)^n, M_{t_0,1})$, while the use of homotopy maps defines a A_∞ functor to $\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q}_{t_0,1})$ which one can easily see is an A_∞ equivalence by an energy filtration argument. Leaving the details to the reader, we conclude

Theorem 5.8. *The categories $\mathcal{T}\text{Fuk}((\mathbb{C}^\star)^n, M_{t_0,1})$ and $\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q}_{t_0,1})$ are equivalent as A_∞ -pre-categories.* \square

We now return to the practice of dropping the subscripts from $M_{t,1}$ and all associated objects.

6. HMS FOR TORIC VARIETIES: MORSE = ČECH

In this section, we shall prove the equivalence of the category $\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$ with a DG category whose objects are holomorphic line bundles on X , and whose morphisms are Čech complexes. We will use the following version of Čech theory:

Given a sheaf of vector spaces (or free abelian groups) \mathcal{S} , and an ordered cover $\{O_k\}_{k=0}^N$ of a (nice) space, such that \mathcal{S} is acyclic on all iterated intersections of the elements of our cover, we consider the Čech complex

$$\check{C}^i(\mathcal{S}) \equiv \bigoplus_{k_0 < \dots < k_i} \Gamma(O_{k_0} \cap \dots \cap O_{k_i}, \mathcal{S}),$$

with the differential on $\Gamma(O_{k_0} \cap \dots \cap O_{k_i}, \mathcal{S})$ defined by

$$(6.1) \quad d\phi = \sum_{\substack{k \\ k_0 < \dots < k_j < k < k_{j+1} < \dots < k_i}} (-1)^{j+1} \phi|_{O_{k_0} \cap \dots \cap O_{k_j} \cap O_k \cap O_{k_{j+1}} \cap \dots \cap O_{k_i}}.$$

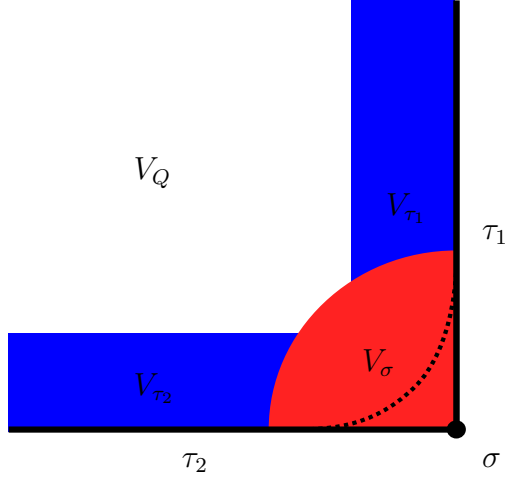


FIGURE 7. The cover V_τ in dimension 2. The dashed line represents the boundary of $\partial\mathcal{Q}$

Further, we define a cup product

$$\check{C}^i(\mathcal{S}_1) \otimes \check{C}^j(\mathcal{S}_2) \rightarrow \check{C}^{i+j}(\mathcal{S}_1 \otimes \mathcal{S}_2)$$

which, on $\Gamma(O_{k_0} \cap \dots \cap O_{k_i}, \mathcal{S}_1) \otimes \Gamma(O_{l_0} \cap \dots \cap O_{l_j}, \mathcal{S}_2)$ is defined by

$$(6.2) \quad \phi \otimes \psi \mapsto \begin{cases} \phi \otimes \psi|_{O_{k_0} \cap \dots \cap O_{k_i} \cap O_{l_1} \cap \dots \cap O_{l_j}} & \text{if } k_i = l_0 \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to check that the above cup product is associative. Since our covers will be labeled by the cells of Q , we will assume that we have chosen an ordering of these cells.

6.1. Another cover of the moment polytope. We inductively construct a cover of the moment polytope analogous to the one used in [1, Section 5.2]. We will write $B_c(Z)$ for (closed) balls of radius c about a subset Z .

Consider a sequence of strictly positive real numbers $\epsilon_{n-1} < \dots < \epsilon_0$, and inductively define V_τ to be the closure of

$$B_{\epsilon_k}(\tau) - \bigcup_{\sigma \subset \partial\tau} \text{int}(V_\sigma),$$

for τ a cell of dimension k . For intuition, the reader should consult Figure 7.

Proposition 6.1. *If $\epsilon_{n-1} \ll \dots \ll \epsilon_0 \ll 1$, then the cells V_τ define a cellular subdivision of Q which is dual to the barycentric subdivision.*

Sketch of proof. It is clear that the intersections of neighbourhoods of cells of the same dimension are localized near their common boundaries. In particular, picking $\epsilon_k \ll \epsilon_{k-1}$ guarantees that the cells V_τ do not have overlapping interior (the construction only guarantees this if the cells of Q have different dimension). Once this condition is satisfied, there is a bijective correspondence between the maximal dimensional cells of the cellular subdivision and the original cells of Q , so it suffices to understand when two of these cells intersect. In particular, we must prove

that the $n - k$ -dimensional cells of the cellular decomposition we constructed are in bijective correspondence with iterated inclusions of cells $\sigma_0 \subset \dots \subset \sigma_k \subset Q$.

We illustrate the case $k = 1$; no new ideas are need for the higher dimensional cases. Assume that σ and τ are respectively of dimension i and j . Note that if σ and τ do not have a common boundary cell, the distance between them is so large that any reasonable choice of ϵ_0 guarantees that no intersection occurs. If they do share a boundary cell ρ of dimension k , then the intersection of ϵ_i and ϵ_j neighbourhoods of σ and τ occur in a neighbourhood of ρ . Choosing ϵ_i and ϵ_j much smaller than ϵ_k therefore guarantees that the boundaries of the cells V_σ and V_τ do not intersect unless, possibly, if $\sigma \subset \tau$.

We must now prove that all such intersections do, in fact, occur. This is clear if σ has codimension 1 in τ . By taking a normal slice, we reduce the problem to the case where σ is a vertex and $\tau = Q$. So we must show that

$$S_{\epsilon_0}(\sigma) - \bigcup_{\sigma \neq \tau' \subset \partial Q} V_{\tau'}$$

is a non-empty cell, where $S_{\epsilon_0}(\sigma)$ is the sphere of radius ϵ_0 in Q . Again, this can be clearly achieved by picking constants $\epsilon_{n-1} \ll \dots \ll \epsilon_0$. Indeed, for such a choice of constants, the $n - 1$ -dimensional volume of the $S_{\epsilon_0}(\sigma)$ is much larger than the volume of its intersection with the cells $V_{\tau'}$. \square

Remark 6.2. The specific choice of coefficients depends on geometry of the polytope Q . Even if we restrict to polytopes corresponding to smooth toric varieties, the “angles” at vertices can still be arbitrarily acute. This, of course, is a result of the fact that, by using a metric, we are restricting the symmetries to $O(n, \mathbb{Z})$, whereas constructions in tropical geometry are usually invariant under the larger group $GL(n, \mathbb{Z})$.

Corollary 6.3. *If, for each cell $\tau \in Q$, W_τ is a sufficiently small open neighbourhood of V_τ , then non-empty intersections*

$$W_{\sigma_0} \cap \dots \cap W_{\sigma_k}$$

correspond to nested inclusions $\sigma_0 \subset \dots \subset \sigma_k$. \square

We observe that if we pick \mathcal{Q} arbitrarily close to Q , we can ensure that the above cover yields a cover of \mathcal{Q} . Further, the construction of \mathcal{Q} and of tropical Lagrangian sections relied on a choice of a cover O_τ of \mathcal{Q} .

Lemma 6.4. *If the parameters which appear in the construction of tropical Lagrangian sections are chosen to be sufficiently small, then*

$$O_\tau \subset \bigcup_{\sigma \subset \bar{\tau}} V_\sigma$$

\square

In particular, the cells $\mathcal{Q} \cap V_\tau$ define a cellular subdivision of \mathcal{Q} , which we denote $\tilde{\mathcal{Q}}_b$. Abusing notation, we will still refer to its maximal cells as V_τ .

6.2. Tropical Lagrangians and simplicial cochains. In order to use the theory developed in earlier sections to relate (relative) simplicial cochains to Morse cochains, we must know that the gradient flow of functions defining tropical Lagrangian sections is well-behaved at the boundary. The technical result that we need is:

Lemma 6.5. *If f is a smooth function on \mathcal{Q} defining the lift of a tropical Lagrangian section whose gradient vector field is transverse to the cells of $\check{\mathcal{Q}}_b$, then, for each cell τ , either $\partial_f^+ \mathcal{Q} \cap V_\tau$ is empty, or its inclusion in V_τ is a deformation retract.*

Proof. It suffices to prove that $\partial_f^+ \mathcal{Q} \cap V_\tau$ is contractible when it is not empty. Our transversality assumption implies that we may restrict our attention to maximal cells V_τ .

We first prove that $\partial_f^+ \mathcal{Q} \cap O_\tau$ is contractible for each cell τ . On this intersection, $\partial \mathcal{Q}$ splits as a product with one factor τ and, in $\partial \mathcal{Q} \cap O_\tau$, the gradient flow of f is only allowed to vary in the direction of τ . By taking a normal slice, the problem is reduced to the case where τ is a vertex, and the gradient flow of f is constant. Let us write v for the corresponding vector in \mathbb{R}^n , and $\partial_v^+ Q$ for the (closure) of the set of maximal faces of Q on which the vector v points outwards.

By construction, $\partial \mathcal{Q}$ is a smoothing of the convex polytope Q , and is itself convex. Observe that $O_\tau \cap \partial \mathcal{Q}$ is a convex subset of $\partial \mathcal{Q}$ (it was constructed using cutoff functions with radial support). Assume (for simplicity) that $\partial \mathcal{Q}$ is strictly convex. In this case, since the Gauss map to S^n has no critical points and $\partial_f^+ \mathcal{Q}$ is the inverse image of a hemisphere, the subset $O_\tau \cap \partial \mathcal{Q}$ is contractible. In our situation, the Gauss map is not a diffeomorphism, but can be easily seen to preserve contractibility.

To prove that $\partial_f^+ \mathcal{Q} \cap V_\tau$ is contractible, we note that for each σ with $\tau \subset \partial \sigma$, we have a non-empty intersection $V_\tau \cap O_\sigma$. We can apply the same analysis to these various intersections to conclude that $\partial_f^+ \mathcal{Q} \cap V_\tau \cap O_\sigma$ is again contractible. Therefore $\partial_f^+ \mathcal{Q} \cap V_\tau$ is covered by contractible sets, so it suffices to prove that the geometric realization of this cover is contractible. However, $\partial_f^+ \mathcal{Q} \cap V_\tau \cap O_\sigma$ is C^0 close to $\partial_v^+ Q \cap V_\tau \cap O_\sigma$ so the combinatorics of this cover of $\partial_f^+ \mathcal{Q} \cap V_\tau$ are equivalent to those of the cover of $\partial_v^+ Q \cap V_\tau$ by the intersections $\partial_v^+ Q \cap V_\tau \cap O_\sigma$. Since $\partial_v^+ Q \cap V_\tau$ is easily seen to be contractible, the result follows. \square

We define $\partial_f^+ \check{\mathcal{Q}}_b$ to be the union of the cells of $\check{\mathcal{Q}}_b$ which intersect $\partial_f^+ \mathcal{Q}$ (this is not a subcomplex, but rather the complement of one). We write $\partial_f^+ Q_b$ for the dual subcomplex of Q_b , and $\partial_f^+ Q$ for the cells of Q corresponding to vertices of $\partial_f^+ Q_b$. As a direct consequence of the above analysis, we have

Corollary 6.6. *If τ is a maximal cell of Q , then $\tau \subset \partial_f^+ Q$ if and only if there is a vertex σ of $\partial \tau$ such that $\text{grad}(f)(\sigma)$ is outwards pointing on τ . Further, $\partial_f^+ Q$ is the closure of its maximal cells.* \square

Let

$$H: \partial Q \rightarrow \mathbb{Z}$$

be a function which is constant on each maximal cell, (the function is likely to be discontinuous on the $n-2$ skeleton). One natural way to obtain such a function H is to consider a lattice vector $u \in \mathbb{Z}^n$. Indeed, since Q is an integral polytope, we have (integral) normal vectors ν_τ which allows us to define a function

$$\begin{aligned} u: \partial Q &\rightarrow \mathbb{Z} \\ u|_\tau &= \langle u, \nu_\tau \rangle. \end{aligned}$$

As in Appendix E, we set

$$\partial_H^+ Q$$

to be the closure of the maximal cells where H is positive. For each such function, we can define a sheaf $\mathcal{S}(H)$ on Q by

$$\Gamma(U, \mathcal{S}(H)) = \begin{cases} 0 & \text{if } U \cap \partial_H^+ Q \neq \emptyset \\ \mathbb{C} & \text{otherwise.} \end{cases}$$

Using the ordering of the cells of Q , we write $W_k = W_{\sigma_k}$, we observe that

$$\Gamma(W_{k_0} \cap \cdots \cap W_{k_i}, \mathcal{S}(H)) = \begin{cases} 0 & \text{if } W_{k_j} \in \partial_H^+ Q \text{ for all } 0 \leq j \leq k_i \\ \mathbb{C} & \text{otherwise.} \end{cases}$$

The fact that

$$\partial_{H_1+H_2}^+ Q \subset \partial_{H_1}^+ Q \cup \partial_{H_2}^+ Q,$$

implies that there is a map of sheaves

$$(6.3) \quad \mathcal{S}(H_1) \otimes \mathcal{S}(H_2) \rightarrow \mathcal{S}(H_1 + H_2).$$

We consider the DG pre-category

$$\check{\text{Cech}}^{\mathbb{Z}^n}(Q)$$

whose objects are such integral valued function on ∂Q and where morphisms between H_0 and H_1 are given by

$$\bigoplus_{u \in \mathbb{Z}^n} \check{C}^*(\mathcal{S}(H_1 - H_0 + u)),$$

and multiplication is given by the composition of the cup product with the map on Čech complex induced by Equation (6.3).

Proposition 6.7. $\check{\text{Cech}}^{\mathbb{Z}^n}(Q)$ and $\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$ are equivalent as A_∞ pre-categories.

Proof. We pick a normal vector field $\nu_{\partial \mathcal{Q}}$ on $\partial \mathcal{Q}$ which agrees with the integral normal vector ν_τ on every maximal cell τ on $O_\tau \cap \partial \mathcal{Q}$. Starting with a function f defining a tropical Lagrangian section, we consider two functions

$$H(f): \partial \mathcal{Q} \rightarrow \mathbb{R} \quad \hat{H}(f): \partial Q \rightarrow \mathbb{Z}.$$

The first function is defined as in Appendix E to take value $\langle \text{grad}(f)(p), \nu_{\partial \mathcal{Q}} \rangle$ at a point p . The second function, on every maximal cell τ , is given by

$$\langle \text{grad}(f)(p), \nu_\tau \rangle$$

for any $p \in O_\sigma \cap \partial \mathcal{Q}$ with $\sigma \subset \partial \tau$. The function $\hat{H}(f)$ is well defined since the restriction of the gradient vector field of f to $O_\sigma \cap \partial \mathcal{Q}$ varies only in directions tangent to σ , hence the inner product with the normal vector of τ does not depend on p or σ . By taking σ to be a vertex, we see that $\hat{H}(f)$ is, indeed, integral valued.

To define the functor, we consider an intermediate category $\text{Simp}^{\mathbb{Z}^n}(Q_b)$ whose objects are the functions $H(f)$, for f an object of $\text{Morse}^{\mathbb{Z}^n}(Q)$, and where morphisms between H_0 and H_1 are given by relative simplicial cochains

$$\bigoplus_{u \in \mathbb{Z}^n} C^*(Q_b, \partial_{H_1 - H_0 + u}^+ Q_b).$$

Claim. $\check{\text{Cech}}^{\mathbb{Z}^n}(Q)$ and $\text{Simp}^{\mathbb{Z}^n}(Q_b)$ are equivalent DG categories.

Proof. The functor take $H(f)$ to $\hat{H}(f)$. We claim that there is a natural isomorphism between the chain complexes on either side. Indeed, in our construction of V_τ and W_τ , we already observed that there is a bijective correspondence between k -cells of the barycentric subdivision and intersections of $k + 1$ open elements of the cover $\{W_\tau\}_{\tau \subset Q}$. It is easy to check that taking cochains relative cells in $\partial_{H(f_1)-H(f_0)+u}^+ Q_b$ exactly corresponds to taking coefficients in the sheaf $\mathcal{S}(\hat{H}(f_1) - \hat{H}(f_0) + u)$, since this sheaf is trivial on open sets which intersect $\partial_{H_1-H_0+u}^+ Q_b$. Both have the effect of excluding generators corresponding to iterated inclusions of cells where all of the cells lie in $\partial_{H_1-H_0+u}^+ Q_b$. That the differential and the composition agree can be seen by the explicit formulae (e.g, Equation (6.2) is just rewriting Equation (E.2) in the dual basis). \square

Having proved the claim, it now suffices to prove that $\text{Simp}^{\mathbb{Z}^n}(Q_b)$ and $\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$ are equivalent. This is done by interpreting Q_b as the triangulation dual to $\check{\mathcal{Q}}_b$. We are then in the context where we can apply Corollary 4.33. There are two issues to consider

- (1) We're working with functions that have specific behaviour at the boundary (i.e. the tropical condition and the integrality condition on the functions H away from the $n - 2$ skeleton). These two conditions simply restrict the set of objects we're considering, without affecting the morphisms. In other words we consider subcategories of $\text{Simp}(Q_b)$ and $\text{Morse}(\mathcal{Q})$. An equivalence of categories as in Corollary 4.33 restricts to an equivalence of categories between every subcategory and its image.
- (2) The morphisms in the categories $\text{Simp}^{\mathbb{Z}^n}(Q_b)$ and $\text{Morse}^{\mathbb{Z}^n}(\mathcal{Q})$ are \mathbb{Z}^n graded. However, by construction, the differential and the composition respect this grading.

\square

6.3. The Čech category of line bundles. In this section, we complete the proof of the main theorem by constructing an equivalence from $\check{\text{Čech}}^{\mathbb{Z}^n}(Q)$ to the Čech category of line bundles. We're taking the dual point of view to that of [18, Chapter 3]. Let X be a projective toric variety, and, as before, let Q denote its moment polytope. The i -cells of Q correspond to invariant, irreducible, i dimensional toric subvarieties of X .

In order to have a chain-level model for these cohomology groups, we will use the Čech complexes associated to the canonical cover by X by affine toric varieties. Equivalently, the polyhedron Q has a natural cell division to which we can associate a cover by open stars of cells. Define

$$U_i = U_{\sigma_i} \equiv \bigcup_{\tau | \sigma_i \subset \bar{\tau}} \tau.$$

This gives a open cover of Q , but, considering the inverse image of cells under the moment map, we also obtain a cover of X . Basic techniques in toric geometry, see [18, Chapter 3], imply that this is an affine cover (the affine toric variety corresponding to a cell σ_i is defined by the normal cone of σ_i).

Let L be a line bundle on X , and let us choose an (algebraic) trivialization over the open dense torus contained in X . Since algebraic functions on $(\mathbb{C}^*)^n$ correspond to Laurent polynomials, we can identify the lattice \mathbb{Z}^n with a basis of sections of $L|_{(\mathbb{C}^*)^n}$. Such sections extend to meromorphic sections of L over X . In particular,

for each Laurent polynomial u , we can record the order of the pole or zero of the extension to every toric divisor. We take the negative of the order, and obtain a \mathbb{Z} valued function on the facets of Q , which we extend to a (not necessarily continuous) function $H(L, u): \partial Q \rightarrow \mathbb{Z}$. Let us write $\partial_{H(L, u)}^+ Q$ for the closure of the set where $H(L, u)$ is positive. This is the set of poles of the extension of u to the toric variety X . The Čech complex $\check{C}^*(L)$ splits as a direct sum labeled by lattice vectors $u \in \mathbb{Z}^n$ with u -factor

$$(6.4) \quad \check{C}_u^d(L) \equiv \bigoplus_{\substack{j_0 < j_1 < \dots < j_d \\ \sigma_{j_k} \notin \partial_{H(L, u)}^+ Q \text{ for all } 0 \leq k \leq d}} \Gamma(U_{j_0} \cap \dots \cap U_{j_d}, \mathbb{C}),$$

where $\Gamma(U_{j_0} \cap \dots \cap U_{j_d}, \mathbb{C})$ consists of all constant maps from $U_{j_0} \cap \dots \cap U_{j_d}$ to \mathbb{C} . Since in fact all these intersections are connected, we simply obtain a copy of \mathbb{C} for each such intersection.

To obtain a category, we set

$$\check{C}_u^d(L_0, L_1) \equiv \check{C}_u^d(L_1 \otimes L_0^{-1}),$$

and define the cup product as follows: if $\phi \in \Gamma(U_{j_0} \cap \dots \cap U_{j_d}, \mathbb{C})$, and $\psi \in \Gamma(U_{i_0} \cap \dots \cap U_{i_{d'}}, \mathbb{C})$, set

$$(6.5) \quad \phi \cup \psi = \begin{cases} \phi \cdot \psi|_{U_{j_1} \cap \dots \cap U_{j_d} \cap U_{i_1} \cap \dots \cap U_{i_{d'}}} & \text{if } j_d = i_0 \\ 0 & \text{otherwise.} \end{cases}$$

The above formula can be interpreted as a cup product

$$\check{C}_{u'}^{d'}(L_1, L_2) \otimes \check{C}_u^d(L_0, L_1) \rightarrow \otimes \check{C}_{u+u'}^{d+d'}(L_0, L_2).$$

Definition 6.8. *The DG category*

$$\check{\text{Cech}}(X)$$

has objects line bundles on X , morphisms the Čech complexes with respect to the affine toric cover of X defined by Equation (6.4), and composition given by the product on Čech complexes defined by Equation (6.5).

Proposition 6.9. *There is an equivalence of categories*

$$\check{\text{Cech}}(X) \mapsto \check{\text{Cech}}^{\mathbb{Z}^n}(Q)$$

Proof. At the level of objects, the functor takes a line bundle L to the function $H(L, u)$. Scrutinizing the morphisms, we observe that Equation (6.4) is simply a low-tech method of recording the Čech complex of $\mathcal{S}(H(L, u))$ with respect to the cover U_i . Since both covers are acyclic, the inclusion $W_i \subset U_i$ induces an equivalence of categories. \square

APPENDIX A. A SECOND LOOK AT SEMI-TROPICAL DEGENERATIONS

In this Appendix, we will modify the construction of semi-tropical degenerations and admissible Lagrangians in order to prove Proposition 5.5. First, we will prove a sharper convergence result than the one used in [1], then we will discuss the three parts of Proposition 5.5

A.1. Choosing $M_{t,1}$ C^k -close to M_t . Fix the partition O_τ of \mathbb{R}^n , and a time t_0 for which we can carry through the construction of $M_{t_0,s}$. We will write $O_\tau(t)$ for $\log(t) \cdot O_\tau$. The following subsets of Q will be useful in what follows:

$$\begin{aligned} O_{\partial Q}(t) &= \bigcup_{\tau \in \partial Q} O_\tau(t) \\ O_k(t) &= \bigcup_{\substack{\tau \in \partial Q \\ \dim(\tau) \leq k}} O_\tau(t) \\ O_Q(t) &= \log(t) \cdot Q - O_{\partial Q}(t). \end{aligned}$$

We will also write $O_\tau^{\mathbb{C}^*}(t) \subset (\mathbb{C}^*)^n$ for the inverse image of $O_\tau(t)$ under the logarithmic moment map, and by extension the meaning of $O_k^{\mathbb{C}^*}(t)$ should be clear.

While the distance between $\partial \mathcal{Q}_{t,s}$ and Q grows logarithmically, we can in fact choose the semi-tropical amoeba degeneration such that, for any s , the distance between $\partial \mathcal{Q}_{t,0}$ and $\partial \mathcal{Q}_{t,s}$ decreases with t .

Proposition A.1. *For any integer k , and each ϵ sufficiently small, there is an appropriate choice of parameters such that for sufficiently large t , the C^k distance between M_t and $M_{t,s}$ is proportional to $t^{-\epsilon}$.* \square

Here, the C^k norm is taken with the respect to the metric on $(\mathbb{C}^*)^n$ induced by the flat metric on \mathbb{R}^n . The proof of the above proposition follows from elementary estimates of the type performed in [1, Section 3]. We explain the main ideas:

First, recall that tropical geometry produces a polyhedra complex $\Pi_\nu \subset \mathbb{R}^n$ (tropical amoeba) from a Laurent polynomial W and a map $\nu: A \rightarrow \mathbb{R}$, where $A \subset \mathbb{Z}^n$ is the subset of monomials which have non-zero coefficients. For simplicity, we assume that the origin lies in A with coefficient -1 , that ν vanishes at the origin, and that all other coefficients are equal to 1. In particular, we shall write

$$(A.1) \quad W_t = -1 + \sum t^{-\nu(\alpha)} z^\alpha.$$

We assume that ν induces a maximal subdivision, in particular the components of the complement of Π_ν are labeled by the elements of A , and each monomial is represented. Concretely, consider the function

$$u \mapsto \langle u, \alpha \rangle - \nu(\alpha)$$

on \mathbb{R}^n . The component C_α is the one where the function corresponding to α dominates all others. The most elementary result relating tropical and complex geometry is:

Lemma A.2. *If \mathcal{A}_t is the amoeba of $M_t = W_t^{-1}(0)$, then, as $t \rightarrow +\infty$, the rescaled amoeba $\frac{1}{\log(t)} \mathcal{A}_t$ converges in the Hausdorff topology to Π_ν .* \square

The idea of the proof (and the proof itself) is extremely simple. Consider a point which is $\epsilon \log(t)$ away from $\log(t) \Pi_\nu$. It must lie in one of the components $\log(t) C_\alpha$. Using the fact that in the coordinates coming from the identification of $(\mathbb{C}^*)^n$ with $T^*\mathbb{R}^n/\mathbb{Z}^n$, we have $z^\alpha = e^{\langle \alpha, u + i\theta \rangle}$, we easily compute that the monomial of W_t labeled by α dominates all others by a factor of $e^{\log(t)\epsilon}$. If t is large enough, the remaining terms are too small for the sum to vanish. In particular, no element of $W_t^{-1}(0)$ can project to a point whose distance to $\log(t) \Pi_\nu$ is larger than $\epsilon \log(t)$.

To see to what extent this estimate is optimal, we consider the example where

$$W_t = -1 + t^{-1}z_1 + t^{-1}z_2.$$

The tropical amoeba has a vertex at $(1, 1)$ at which horizontal, vertical, and diagonal infinite edges meet. In particular, $(\log(t), \log(t))$ is the vertex of the rescaling. Note that the point $(\log(t/2), \log(t/2))$ lies on the amoeba \mathcal{A}_t . The distance between this point and $\log(t)\Pi_\nu$ is independent of t and equal to $\log(\frac{1}{2})$; in particular, the amoeba \mathcal{A}_t does not converge to Π_ν . However, if $M_{t,\nu}$ is the complex tropical hypersurface lying over $\log(t) \cdot \Pi_\nu$ then it is easy to prove a general result:

Lemma A.3. *If O_{n-2} contains the ϵ neighbourhood of the $n-2$ skeleton of Π_ν , then the distance in the complement of $O_{n-2}(t)$ between the amoeba \mathcal{A}_t and $\log(t)\Pi_\nu$ is bounded in norm by a constant multiple of $t^{-\epsilon}$. Moreover, the distance in the C^k topology between corresponding submanifolds of $(\mathbb{C}^*)^n - O_{n-2}^*(t)$ is also proportional to $t^{-\epsilon}$. \square*

Again, the proof is completely elementary. To simplify the explanation, we will assume that we are working near $\log(t)C_0$, the region where the constant term dominates. This is the only situation we care about in the desired application.

We illustrate C^0 convergence. Let us restrict our attention to an $n-1$ cell of $\log(t)\Pi_\nu$ where the function defined by some monomials α has norm 1. Consider a point outside the δ -neighbourhood of this cell, say in the direction of α . We know that, at such a point, the monomial corresponding to α has norm larger than e^δ . In particular, the difference

$$-1 + t^{-\nu(\alpha)}z^\alpha$$

has norm larger than δ . Now, the norm of the remaining monomials is essentially bounded by $t^{-\epsilon}$ away from the $\epsilon \log(t)$ neighbourhood of the $n-2$ -skeleton. This means that if we pick δ to be $Kt^{-\epsilon}$ for some fixed constant K , these remaining monomials are not large enough for the sum to vanish. In particular, every point of \mathcal{A}_t which is more than $\epsilon \log(t)$ away from the $n-2$ -skeleton of $\log(t)\Pi_\nu$ lies within $Kt^{-\epsilon}$ of the $n-1$ -skeleton. In order to prove C^0 -closeness in $(\mathbb{C}^*)^n$, it suffices to take into account the norm of both the real and imaginary parts of $W_{t,0}$.

The same analysis yields C^k bounds. Indeed, we're using a flat metric with respect to the (u, θ) variables, so if we write $z^\alpha = e^{\langle \alpha, u + i\theta \rangle}$, we compute that the action of dz^α on tangent vectors is given by

$$(A.2) \quad dz^\alpha(-) = \langle \alpha, - \rangle e^{\langle \alpha, u + i\theta \rangle}$$

In particular, the norm of dz^α is equal to $|\alpha| \cdot |z^\alpha|$, so the bounds on C^0 norms of each monomial imply the bounds on C^1 norms. We conclude that away from the $\log(t)\epsilon$ neighbourhood of the $n-2$ skeleton, the leading order term of dW_t has norm which is on the order of t^ϵ larger than the norm of any other term. This proves C^1 -closeness, and the reader may easily check that the same estimate holds for higher derivatives.

To extend the analysis beyond the maximal cells, we consider a cell $\sigma \subset \Pi_\nu$ of arbitrary dimension, and the polynomial W_t^σ consisting of the monomials in W_t corresponding to the maximal cells which “meet at σ .” We write M_t^σ for the zero level set of this function. Assuming that O_τ contains an ϵ neighbourhood of τ , for every cell τ , the same proof technique yields the following result:

Lemma A.4. *In the C^k norm, the distance between M_t and M_t^σ is proportional to $t^{-\epsilon}$ on $O_\sigma^{\mathbb{C}^*}(t)$ (the subset of $(\mathbb{C}^*)^n$ which projects (under the logarithmic moment map) to $O_\sigma(t)$). \square*

The proof of Proposition A.1 should now be clear, since the function $W_{t,1}$ is essentially patching the various polynomials W_t^σ over the skeleton of Π_ν . As a consequence of the fact that the C^k -norm of W_t^σ is bounded by the C^0 -norm of its monomial terms, we conclude

Corollary A.5. *The C^k -norms of W_t and $W_{t,1}$ are uniformly bounded in $\log(t)C_0$ (the neighbourhood where the constant term dominates).*

A.2. Construction a C^k -small symplectomorphism. The next step of the construction of the category of tropical Lagrangian sections involved the choice of a symplectomorphism $\phi_{t,1}$ of $(\mathbb{C}^*)^n$ mapping M_t to $M_{t,1}$. We assume ϵ has been fixed as in Proposition A.1.

Proposition A.6. *For any integer k , and δ sufficiently small, there is an appropriate choice of parameters such that for sufficiently large t , the symplectomorphism*

$$\phi_{t,1}: ((\mathbb{C}^*)^n, M_t) \rightarrow ((\mathbb{C}^*)^n, M_{t,1})$$

is supported in a neighbourhood of M_t of size δ and has C^k norm controlled by $\frac{t^{-\epsilon}}{\delta^k}$. Moreover, we can assume that $\phi_{t,1}$ preserves the positive real locus in $(\mathbb{C}^)^n$. \square*

To produce the desired symplectomorphism, we appeal to Weinstein's theorem for symplectic submanifolds "with bounds." First, we prove a bound on the size of the Weinstein neighbourhoods of M_t :

Lemma A.7. *There exists a constant δ_0 such that, for sufficiently large t , the δ_0 neighbourhood of M_t is embedded and is contained in a Weinstein neighbourhood of M_t . Moreover, for t sufficiently large, the curvature of M_t is uniformly bounded.*

Proof. By the previous section, the hypersurfaces M_t and M_t^σ are uniformly C^k -close in $O_\sigma^{\mathbb{C}^*}(t)$ so it suffices to prove the result for each M_t^σ . Note that this is the product of a flat hyperplane with the zero level set of a general hyperplane in a complex torus of dimension $\text{codim}(\sigma)$. We can always write this family as the zero level set of

$$(A.3) \quad -1 + \sum_{i=0}^k t^{-\nu_i} z^{\alpha_i}$$

where the vectors α_i are linearly independent and form a basis for the integral subspace of \mathbb{Z}^n that they span. Let A_σ denote the matrix whose rows are α_i , and $\vec{\nu}_\sigma$ the column vector whose entries are ν_i . It is easy to check that translation (in the base \mathbb{R}^n) by any vector \vec{b} which is solution to

$$\vec{b} \cdot A_\sigma = b \vec{\nu}_\sigma$$

intertwines the hypersurfaces M_t^σ and M_{t+b}^σ , and hence their Weinstein neighbourhoods. The desired result follows. \square

The proof of Proposition A.6 will now follow from general considerations. Let Y be a symplectic manifold equipped with a Riemannian metric, and M a symplectic submanifold. By Weinstein's results, there exists a neighbourhood of M which is symplectomorphic to a neighbourhood of the zero section in the total space E of

the normal bundle of M . Choose $\delta < 1$ such that the δ -neighbourhood of M is contained in such a Weinstein neighbourhood.

Lemma A.8. *If M_s is a family of symplectic submanifolds for $s \in [0, 1]$ with $M_0 = M$ whose distance from M in the C^k topology (for $k > 0$) is bounded by ϵ , and $\epsilon < \delta$ then there exists a symplectomorphism*

$$\phi_s: (Y, M) \rightarrow (Y, M_s)$$

satisfying the following properties

- (1) *The C^k norm of ϕ_s is bounded by $C \frac{\epsilon}{\delta^k}$ with proportionality constant depending only on the curvatures of Y , M , and ω .*
- (2) *The support of ϕ_s is contained in a δ neighbourhood of M_0 .*

Sketch of proof: The assumptions of C^k -closeness imply that M_s lies in the δ neighbourhood of M , and is moreover a graph of a C^k -small section of the normal bundle under the identification from Weinstein's theorem. Choosing our constant C sufficiently large, we can perform all estimates in the normal bundle. Note that the restriction of the symplectic form to each M_s is C^k -close to the pullback of the symplectic form from M . Further, the two forms are clearly cohomologous. In particular, applying the usual proof of Moser's stability theorem for symplectic forms, we find, along each M_s , a vector field X_s with values in TY whose C^k norm is controlled by a universal constant times ϵ , and such that the flow of X_s defines a symplectomorphism from M to M_1 .

Using the symplectic form, we turn X_s into a 1-form on M_s valued in T^*Y . We find a family H_s of functions which generate a hamiltonian flow mapping M to M_1 . Near M , the C^{k+1} norm of H_s (and hence the C^k -norm of the flow) is still controlled by ϵ . However, we require that each H_s vanish outside the δ -neighbourhood of M in order to be able to extend the flow to Y . Choosing an appropriate cutoff function we can still ensure that the C^k norm is controlled by $\frac{\epsilon}{\delta^k}$ up to a fixed multiplicative constant. \square

Remark A.9. Note that we have much flexibility in the previous construction. In particular, if L is a Lagrangian submanifold of Y such that the intersections $L_s = L \cap M_s$ are Lagrangian submanifolds whose C^k distance is bounded by ϵ , then we may choose vector fields X_s , hence a symplectomorphism ϕ_1 , which preserve L and still satisfy the desired bounds.

A.3. Controlling Morse flows. We now prove the first part of Proposition 5.5. We first introduce yet another partition of the moment polytope Q whose elements we will denote by P_τ , and whose $\log(t)$ rescaling we will denote by $P_\tau(t)$. We will require the following properties:

- (1) For each cell τ , we have inclusions

$$\bigcup_{\sigma \in \partial\tau} O_\sigma \subset \bigcup_{\sigma \in \partial\tau} P_\sigma \subset \bigcup_{\sigma \in \partial\tau} V_\sigma,$$

- (2) and, if P_k is the union of all P_τ labeled by cells τ of dimension less than or equal to k , then $\mathcal{Q}_{t,1} - P_k(t)$ is convex.

The easiest way to produce such a cover is to use supporting hyperplanes for each cell τ , i.e. hyperplanes which intersect Q precisely at τ . For example, at a vertex v , we can set P_v to be the intersection of Q with the half-space defined by a

hyperplane parallel to a supporting hyperplane at v . Note that O_v can be assumed to be arbitrarily small, so it is easy to achieve the condition

$$O_v \subset P_v \subset V_v.$$

Next, one would define P_τ for an edge τ by taking the intersection of Q with a half-space whose boundary is parallel to a supporting hyperplane at τ (we can define such a hyperplane to be the product of τ with a supporting hyperplane of the projection of Q to the orthogonal complement). Details for the remainder of the (inductive) construction are left to the reader. As before, we will write $P_{\partial Q}(t)$ for the union of all the sets $P_\tau(t)$ for $\tau \subset \partial Q$, and $P_Q(t)$ for its complement in Q .

Next, consider any function f whose restriction to $\partial \mathcal{Q}_{t,1} \cap P_\sigma(t)$ is of the form

$$(A.4) \quad f_\sigma + \langle u, \psi \rangle + \rho$$

where f_σ depends only the directions tangent to σ , ψ is a fixed integral vector, and for which there exists a constant C such that the differential of ρ satisfies

$$(A.5) \quad 0 < \langle d\rho, v \rangle \leq C|v|r_{\partial \mathcal{Q}_{t,1}}$$

whenever the vector v is perpendicular to σ and points towards the interior of Q . In other words, the differential of ρ is positive on the inward pointing cone at σ , and the positivity is controlled by the distance to the boundary.

Lemma A.10. *If the constant C in Equation (A.5) is sufficiently small, then every gradient trajectory of f with endpoints in $P_Q(t)$ lies entirely within $P_Q(t)$.*

Proof. We will prove that no such gradient trajectory can intersect $P_\sigma(t)$ by considering the component of the flow in the plane perpendicular to σ . By induction, we assume that if $k = \dim(\sigma)$ there is no gradient trajectories which moves from the complement of $P_{k-1}(t)$ into $P_{k-1}(t)$ and then exits into $P_Q(t)$.

By C sufficiently small, we mean that $|d\rho| \ll 1$. If ψ is tangent to σ , then Equation (A.5) implies that the gradient flow points away from the boundary so that no gradient trajectory can enter $P_\sigma(t)$ from $P_\tau(t)$ where τ has higher dimension. Assuming the vector ψ is not tangent to σ , the gradient flow of f on P_σ is C^1 close to the gradient flow of $f_\sigma + \langle u, \psi \rangle$, so that the component of $\text{grad}(f)$ in the plane perpendicular to σ is C^1 close to a linear flow. In particular, the convexity of $\mathcal{Q}_{t,1} - P_k(t)$ implies that a gradient flow line which enters $P_\sigma(t)$ either escapes to the boundary, or moves into $P_\tau(t)$ for $\tau \subset \partial \sigma$, at which point we use the inductive assumption to conclude that the gradient flow line cannot return to $P_Q(t)$. \square

If f defines a tropical Lagrangian section, then its restriction to $V_\sigma(t_0)$ can be written as

$$(A.6) \quad f_\sigma^{t_0} + \langle u, \psi \rangle + \rho^{t_0}$$

where ρ vanishes, together with its differential, at the boundary. More precisely, sufficiently close to the boundary, ρ^{t_0} is C^2 -close to a constant multiple of the square of the distance to the boundary. If this constant is positive, then ρ^{t_0} satisfies condition (A.5). In particular, we can deform any function f^{t_0} defining a tropical Lagrangian section so that the term ρ^{t_0} satisfies conditions (A.5) on $P_\sigma(t_0)$ for any sufficiently small constant C , in which case the restriction of $f^{t_0} - v$ to $P_{\partial Q}(t_0)$ has no critical points for any integral covector v . More generally, we have:

Lemma A.11. *If f^{t_0} defines a tropical Lagrangian section over $\mathcal{Q}_{t_0,1}$, and its restriction to $P_\sigma(t_0)$ satisfies Conditions (A.4) and (A.5), then there exists a family of functions f^t defining tropical Lagrangian sections over $\mathcal{Q}_{t,1}$ and which satisfy the following additional properties*

(1) *The restriction of f^t to $P_Q(t)$ is given by*

$$(A.7) \quad f^t(u) = \frac{\log(t)}{\log(t_0)} f\left(\frac{\log(t_0)}{\log(t)} u\right)$$

(2) *The restriction of f^t to $P_\sigma(t)$ can be decomposed as*

$$(A.8) \quad f_\sigma^t + \langle u, \psi \rangle + \rho^t$$

with $d\rho^t$ bounded independently of t , and satisfying the bound (A.5).

Proof. The expression (A.7) defines a function on $\frac{\log(t_0)}{\log(t)} \mathcal{Q}_{t_0,1}$. A look at the construction of $\mathcal{Q}_{t,1}$ reveals that this rescaled copy of $\mathcal{Q}_{t_0,1}$ is contained in $\mathcal{Q}_{t,1}$. In each element $V_\sigma(t)$ of our cover, there is a unique way to extend

$$f_\sigma \left(\frac{\log(t_0)u}{\log(t)} \right)$$

to the complement of $\frac{\log(t_0)}{\log(t)} \mathcal{Q}_{t_0,1}$ preserving the property that the function is independent of the directions perpendicular to σ .

To complete the construction, it suffices to explain how to choose ρ^t . Again, the construction is inductive, so we start with a neighbourhood of a vertex. Start by extending $\rho^{t_0} \left(\frac{\log(t_0)u}{\log(t)} \right)$ to $P_v(t)$ by setting it to vanish outside the intersection with $\frac{\log(t_0)}{\log(t)} \mathcal{Q}_{t_0,1}$; this function is once differentiable. Next, we pick any function ρ' which agrees with the square of the distance to the boundary near $\partial \mathcal{Q}_{t_0,1}$, and satisfies Equation (A.5) for some positive number C . Rescaling ρ' so that it is sufficiently small, we may interpolate between ρ' and $\rho^{t_0} \left(\frac{\log(t_0)u}{\log(t)} \right)$ using cutoff functions so that Equation (A.5) still holds. A C^1 -small perturbation gives a smooth function with the desired property.

If v is a vertex of an edge τ , then the restriction of ρ_v^t to the intersection of $P_v(t)$ with $V_\tau(t)$ satisfies the condition (A.5) for the positive cone of τ . Convexity of this property allows us to extend it to a function ρ_τ^t on the open set $P_\tau(t)$, and the rest of the construction proceeds by induction. \square

Note that our construction of the functions ρ_σ^t does not guarantee a uniform lower bound in Equation (A.5) as we may have to rescale the intermediate function ρ' to ensure that the cutoff still satisfied (A.5). However, we may ensure that the norm of $d\rho^t$ is uniformly bounded in t . Applying Lemma A.10, we conclude the the first part of Proposition 5.5 if we choose δ_1^t such that the δ_1^t neighbourhood of $\partial \mathcal{Q}_{t,1}$ is contained in $P_{\partial Q}(t)$. In the next sections, we will perturb f^t by C^1 -small amounts near $\partial \mathcal{Q}_{t,1}$ to ensure that the remaining parts of Proposition 5.5 also hold; we are free to do this as long as the bound (A.5) holds for some sufficiently small constant C .

A.4. Compactness for $J_{t,s}$ holomorphic curves. We will now modify the family L^t in order to satisfy the second property listed in Proposition 5.5. The idea is to ensure that there is a constant ϵ_t such the images of the Lagrangians L^t agree with a line segment near the circle of radius ϵ_t in \mathbb{C} . We will then prove a compactness

theorem analogous to the one proved in [1, Lemma 2.8], without requiring that the segment extend all the way to the origin.

We start by considering the Laurent polynomials W_t^σ . As in the proof of Lemma A.7, the fact that translation in \mathbb{R}^n intertwines the set of solutions of W_t^σ for different values of t implies that for any constant c_1 , there is a uniform constant 2ϵ independent of σ and t such that the parallel transport of the positive real locus in M_t^σ along the straight-line segment in \mathbb{C} from the origin to $(-2\epsilon, 2c_1\epsilon)$ is a graph over the zero-section, and its defining function satisfies Equation (A.5) for $C = 2c_1$. This latter property holds because we know that the function defining this graph is C^2 -close to a positive multiple of the square of the distance to the boundary, which manifestly satisfies the desired condition in a sufficiently small neighbourhood.

Applying Proposition A.6 and Lemma A.4, which together imply that, as t grows, $M_{t,1}$ becomes C^k -close at every point to some hypersurface M_t^σ we conclude:

Lemma A.12. *For any fixed constants $0 < 2c_1 < C$, there exists a constant ϵ so that the following properties hold for all t sufficiently large:*

- (1) *the support of $\phi_{t,s}$ lies entirely within the inverse image of the ball of radius ϵ under $W_{t,1}$, and*
- (2) *the parallel transport of ∂L^t along a segment from the origin to $(-\epsilon, c_1\epsilon)$ defines a section over a neighbourhood of $\partial \mathcal{Q}_{t,1}$, and*
- (3) *the Lagrangian L^t can be chosen to agree with the above parallel transport near the circle of radius ϵ . Further, the image of L^t under $W_{t,1}$ lies in the left half plane, with all interior intersection points with the zero section mapping to the complement of the disc of radius ϵ .*

Proof. Since the first two parts follow immediately from the fact that as t goes to infinity, the C^k -norm of $\phi_{t,1}$ as well as its support may be assumed to shrink, we only discuss the last part. The proof that the image of L^t may be assumed to lie in the left half plane can be done in two steps. First, we consider the standard hypersurface

$$(A.9) \quad -1 + \sum_i z_i$$

The complement of the amoeba of this standard polynomial has a uniquely determined component whose inverse image in $(\mathbb{C}^*)^n$ is contained in the locus where the norm of all variables z_i is bounded by 1. The boundary of this subset of $(\mathbb{C}^*)^n$ consists of complex numbers z_i whose norms lie on the component of the boundary of the standard amoeba corresponding to the positive real locus; this implies that the image of the boundary under the polynomial (A.9) lies in the left half plane (in fact, it intersects the y -axis only at the origin). Since interior points are obtained by rescaling an appropriate boundary point by a positive real number bounded in norm by 1, and all coefficients in (A.9) are positive, this entire component maps to the left half plane.

Note that by acting by $SL_n(\mathbb{Z})$ and by translations, we conclude that the analogous component of the complement of the amoeba of M_t^σ has inverse image which also maps, under W_t^σ , to the left half-plane. This implies the desired result for $W_{t,1}$ by using C^k -closeness to polynomials W_t^σ in the domains $V_\sigma(t)$ (note that in $V_Q(t)$, $W_{t,1}$ is C^k -close to -1). \square

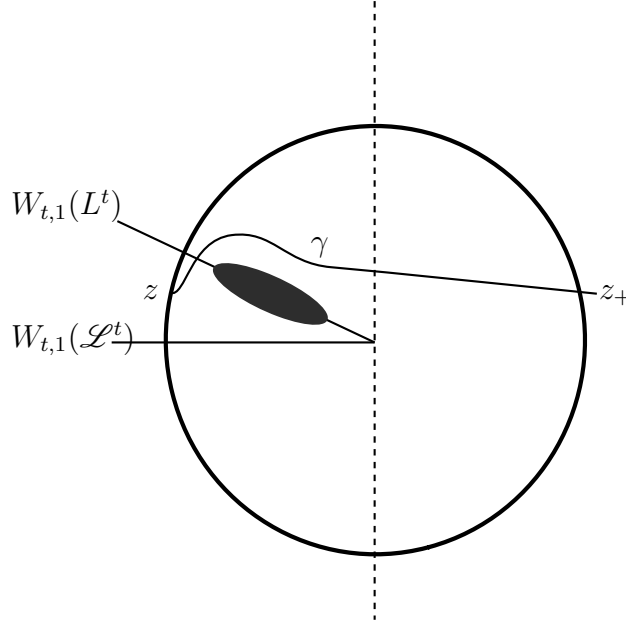


FIGURE 8.

We now prove that choosing L_i^t to agree with the parallel transport of its boundary near a circle of radius ϵ implies a compactness theorem for $J_{t,s}$ holomorphic discs.

Lemma A.13. *If L^t is chosen as in Lemma A.12, then every $J_{t,s}$ holomorphic disc with marked points mapping to interior intersection points of L^t with the zero section projects under $W_{t,1}$ to the complement of the disc of radius ϵ .*

Proof. Assume by contradiction that there exists such a holomorphic disc u which intersects the disc of radius ϵ . Note that $J_{t,s}$ agrees with J near the inverse image of this circle, so that $W_{t,1}$ is $J_{t,s}$ holomorphic near the circle of radius ϵ for every $s \in [0, 1]$. Since all intersection points map outside the circle of radius ϵ , the image of u must intersect the circle of radius ϵ ; let z be such an intersection point whose inverse image under u contains an interior point of the disc. Note that the degree of u at z (number of inverse images counted with signs and multiplicity) is necessarily positive by holomorphicity.

Consider the path γ shown in Figure 8. By applying the maximum principle to the composition of $W_{t,1}$ with the projection to the x -axis, we find that the image of any holomorphic curve under $W_{t,1}$ cannot intersect the complement of the disc of radius ϵ in the right half plane. In particular, the inverse image of z_+ under the composition $W_{t,1} \circ u$ is empty. However, the difference between the degree of u and z and z_+ is given by the intersection number between γ and the image of the boundary of the disc under u . The boundary of the disc consists of a pair of intervals connecting the two intersection points along the zero section and L^t . Since the images of the intersection points lie in the left half plane and away from the disc of radius ϵ , we obtain two paths in the plane, with end points away from the disc of radius ϵ about the origin, and which agree respectively with the negative

x -axis, and the straight line of slope $-c_1$ near the circle of radius ϵ . In particular, the intersection number between γ and the boundary vanishes, which yields the desired contradiction since the degree of u at z cannot vanish. \square

To prove Part (2) of Proposition 5.5 we choose δ_2^t so the inverse image under $W_{t,1}$ of the circle of radius ϵ in L_i^t projects to the complement of the δ_2^t neighbourhood of $\partial\mathcal{Q}_{t,1}$ under the logarithmic moment map. Let us write

$$(A.10) \quad \mathcal{Q}_{t,1}(\delta_2^t)$$

for the complement of this neighbourhood in $\mathcal{Q}_{t,1}$.

The previous result implies that only the interior points of a $J_{t,s}$ -holomorphic curve can map outside to the complement of $\mathcal{Q}_{t,1}(\delta_2^t)$, and that the support of $\phi_{t,s}$ is contained in this complement. In order to exclude this possibility, we recall that the boundary of $\partial\mathcal{Q}_{t,1}$ is convex, so we may choose t large enough and δ_2^t small enough, so that the boundary of $\mathcal{Q}_{t,1}(\delta_2^t)$ is also convex.

The image of any $J_{t,s}$ -holomorphic curve under $\phi_{t,s}^{-1}$ is a J -holomorphic curve, where J is the standard complex structure. Since $\phi_{t,s}$ preserves the boundary of $\mathcal{Q}_{t,1}(\delta_2^t)$, it suffices to prove that no J -holomorphic curve in $(\mathbb{C}^*)^n$ can have a projection which escapes $\mathcal{Q}_{t,1}(\delta_2^t)$. The maximum principle for holomorphic curves in \mathbb{C}^n immediately implies this (See Lemma A.14).

A.5. Compactness for ϵJ holomorphic curves. We finally complete the proof of Proposition 5.5 by proving that ϵJ holomorphic curves with boundary on L^t cannot escape towards the boundary of $\partial\mathcal{Q}_{t,1}$. Our main technique will involve using the product decomposition near every cell τ of ∂Q , then the maximum principle for holomorphic curves with boundary in the factor orthogonal to τ .

Let \mathcal{V} be a convex smooth hypersurface in \mathbb{R}^n , and write $\mathcal{V}^{\mathbb{C}}$ for its inverse image in \mathbb{C}^n under the projection to the real coordinates $\mathbb{C}^n \rightarrow \mathbb{R}^n$.

Lemma A.14. *No holomorphic curve in \mathbb{C}^n with boundary on the zero-section can touch $\mathcal{V}^{\mathbb{C}}$ from the interior.*

Proof. If h is a function on \mathbb{C}^n depending only the real variables whose 0-level set is $\mathcal{V}^{\mathbb{C}}$, then the kernel of the restriction of $d^c h$ to $\mathcal{V}^{\mathbb{C}}$ agrees with the complex distribution. The convexity of \mathcal{V} implies that $d(d^c h)$ is positive on this distribution, so $\mathcal{V}^{\mathbb{C}}$ is J -convex, and we can apply the maximum principle at the interior, see [32]. Since, in addition, $d^c h$ vanishes on $\mathbb{R}^n \subset \mathbb{C}^n$, we have von Neumann boundary conditions, so we can also apply the maximum principle at the boundary. \square

We now return to the problem at hand, and pick any δ_3^t which is smaller than both $\delta_1^t/2$ and $\delta_2^t/2$. The boundary of $\mathcal{Q}_{t,1}$ is convex, so it is easy to pick a convex hypersurface $\mathcal{V} \subset \mathcal{Q}_{t,1}$ satisfying the following properties:

- (1) \mathcal{V} is contained in the $2\delta_3^t$ neighbourhood of $\partial\mathcal{Q}_{t,1}$ but in the complement of the δ_3^t neighbourhood, and
- (2) The intersection of \mathcal{V} with $V_\tau(t)$ splits as an orthogonal product with a factor consisting of directions tangent to τ , and the other given by a convex hypersurface \mathcal{V}_τ .

We pick a function h whose zero level set is \mathcal{V} and for which 0 is a regular value, and write $\alpha = d^c h$ after extending h to \mathbb{C}^n .

Consider the Lagrangian L_0^t , which, on each open set $V_\sigma(t)$ is given as the graph of the differential of the function

$$(A.11) \quad f_\sigma + \langle \psi, - \rangle$$

as in Equation (A.4).

Lemma A.15. *The 1-form $d^c h$ vanishes on L_0^t .*

Proof. It suffices to prove this in each set V_σ where both \mathcal{V} and L_0^t split as Riemannian products. Since L_0^t is flat precisely in the directions where \mathcal{V} is curved the result easily follows. \square

We will now pick ρ^t carefully so that $d^c h$ still vanishes on L^t near its intersection with \mathcal{V} . First, we observe the following

Lemma A.16. *There is a choice of normal vector ν_σ for each face $\sigma \subset \partial Q$ satisfying the following conditions:*

- (1) ν_σ lies in the inward cone at σ , and
- (2) if $\sigma \subset \partial\tau$, then ν_τ agrees with the orthogonal projection of ν_σ along the tangent space of τ .

Proof. Let q_0 be an arbitrary interior vertex of Q , and set ν_σ at a point $p \in \sigma$ to be the projection of the vector from p to q_0 to the orthogonal complement of the tangent space of σ at p . The first property follows from the convexity of Q , and the second from the compatibility of orthogonal projection with restriction to subspaces. \square

We will now require ρ to satisfy the following property in each set $V_\sigma(t)$:

$$(A.12) \quad \text{Near } \mathcal{V}, \text{ the projection of the gradient of } \rho^t \text{ to the plane orthogonal to } \sigma \text{ is constant and agrees with } \nu_\sigma.$$

We are forcing $d\rho^t$ to only vary in direction tangent to σ near the intersection of $V_\sigma(t)$ with \mathcal{V} . One needs to check that there is indeed a function satisfying the above property (e.g. the corresponding 1-form needs to be closed). We can proceed by induction with respect to the dimension of the cell; whenever σ is a vertex, we write ρ_σ for a function whose gradient on $V_\sigma(t) \cap \mathcal{V}$ is fixed to agree with ν_σ . At the end of the construction, ρ_σ shall be the restriction of ρ^t to $V_\sigma(t)$. If σ_1 and σ_2 are the boundary vertices of an edge τ , ρ_{σ_i} defines a function on part of the boundary $V_\tau(t)$ whose restriction to $V_\tau(t) \cap \mathcal{V}$ has normal component which agrees with ν_τ . Consider any extension of ρ_{σ_i} to a function on $V_\tau(t)$ whose gradient has normal component ν_τ . Since the space of such functions is convex, an appropriate linear combination of the functions coming from the two end points of τ gives a function ρ_τ which extends both σ_1 and σ_2 to a function on

$$V_{\sigma_1}(t) \cup V_\tau(t) \cap V_{\sigma_2}(t)$$

satisfying the desired condition. The reader may easily complete the construction for higher dimensional cells, and check that the resulting function ρ^t can be chosen so that Condition (A.5) holds.

The same decomposition argument as in Lemma A.15 implies:

Lemma A.17. *If ρ^t satisfies the condition described in Equation (A.12), then the restriction of $d^c h$ to L_t vanishes near \mathcal{V} . In particular, no J -holomorphic curve with boundary on the zero-section and L_t may be tangent to \mathcal{V} . \square*

Note that by picking ρ^t small enough, we can ensure that the conclusions of the previous two sections still hold. We have therefore completed the proof of the last part of Proposition 5.5.

APPENDIX B. MODULI SPACES OF TREES

In this appendix, we define a stratification into polyhedral cells of certain moduli spaces of trees introduced in Definition 4.25. We begin by noting that the moduli space \mathcal{T}_d admits a decomposition into subsets labeled by the topological type of a tree with $d + 1$ leaves. Let us write $[T]$ for the homeomorphism type of a tree and $\mathcal{T}_d([T])$ for the corresponding subset of \mathcal{T}_d . Note that the lengths of edges defines a homeomorphism

$$(B.1) \quad \mathcal{T}_d([T]) \rightarrow (0, +\infty)^r$$

where r is the number of non-trivial edges. The closure $\tilde{\mathcal{T}}_d([T])$ of $\mathcal{T}_d([T])$ in \mathcal{T}_d is homeomorphic to $[0, +\infty)^r$, with the corner strata corresponding to topological types obtained by collapsing one or more edges of $[T]$. Since every tree can be obtained by collapsing a trivalent tree, we conclude

Lemma B.1. *The moduli space \mathcal{T}_d admits a decomposition into cells which are homeomorphic to $[0, +\infty)^{d-2}$.* \square

Note that the compactification of the moduli space \mathcal{T}_d corresponds to allowing edges to achieve infinite length. We shall denote the closure of each cell $\mathcal{T}_d([T])$ in $\overline{\mathcal{T}}_d$ by $\overline{\mathcal{T}}_d([T])$. Upon fixing an identification of $[0, \infty]$ with $[0, 1]$ (say via the exponential of the negative), we conclude

Lemma B.2. *The compactified moduli space $\overline{\mathcal{T}}_d$ admits a canonical decomposition into subsets $\overline{\mathcal{T}}_d([T])$ which are homeomorphic to cubes.* \square

Note that this decomposition is compatible with the decomposition of the boundary of $\overline{\mathcal{T}}_d$ as the product of lower dimensional moduli spaces. With this in mind, we can now describe the regularity condition on the moduli spaces of gradient trees $\overline{\mathcal{T}}_d$ stated in Definition 4.21.

First, we use the homeomorphisms (B.1) to equip each cell $\tilde{\mathcal{T}}_d([T])$ with a smooth structure as a manifold with corners.

Recall that $\mathcal{T}(q, \vec{p})$ is the moduli spaces of gradient trees with inputs $\vec{p} = (p_1, \dots, p_d)$ and output q . Let us write $\tilde{\mathcal{T}}(q, \vec{p}; [T])$ for those gradient trees whose topological type corresponds to a tree in $\tilde{\mathcal{T}}_d([T])$. In order to prove that this is a smooth manifold, we consider the space

$$\tilde{\mathcal{E}}(q, \vec{p}; [T]) = W^s(q) \times W^u(p_1, \dots, p_d) \times \tilde{\mathcal{T}}_d([T]),$$

where $W^u(p_1, \dots, p_d)$ is the product of the descending manifolds of the critical points $\{p_1, \dots, p_d\}$.

Let us fix a tree $T \in \tilde{\mathcal{T}}_d$, and a basepoint $s \in T$. Let v_j^o be the finite end of each incoming edge e_j . The geodesic A between v_i^o and s determines a map

$$\Phi_j^s: W^u(p_j) \rightarrow Q$$

which is the composition of the gradient flows along the functions labeling the edges of A . The gradient flow is taken for time equal to the length of the corresponding

edge. We can similarly define a map with source $W^s(q)$. Given a choice of basepoints (which is unique up to homotopy since the total space of the tautological bundle over $\tilde{\mathcal{T}}_d$ is contractible), we obtain a smooth map

$$(B.2) \quad \Phi_{[T]}: \tilde{\mathcal{E}}(q, \vec{p}; [T]) \rightarrow Q^{d+1}.$$

The space $\tilde{\mathcal{T}}(q, \vec{p}; [T])$ is inverse image of the diagonal $\Delta^{d+1} \subset Q^{d+1}$ under this map. Sard's theorem implies that this inverse image is a smooth manifold whenever the map $\Phi_{[T]}$ is transverse to the diagonal. The diffeomorphism of Q^{d+1} induced by gradient flow along the edges connecting different basepoints proves that this moduli space and its smooth structure are independent of the choice of basepoints.

Note that the moduli space $\mathcal{T}(q, \vec{p})$ is the union of the spaces $\tilde{\mathcal{T}}_d(q, \vec{p}; [T])$ over their common boundary strata.

Definition B.3. *We say that $\mathcal{T}(q, \vec{p})$ is **regular** if all evaluation maps $\Phi_{[T]}$ are transverse to the diagonal.*

The proof of the next Lemma, which is omitted, is an exercise in Morse theory. The key point is that whenever $\Phi_{[T]}$ is transverse for a fixed metric tree T , it is also transverse for nearby metric trees which can be obtained by resolving vertices of T of valency greater than 2.

Lemma B.4. *For generic choice of functions \vec{f} the moduli space $\mathcal{T}(\vec{f})$ is regular, and admits a compactification $\overline{\mathcal{T}}(\vec{f})$ into a compact manifold with boundary, admitting a decomposition into (smooth) compact manifolds with boundary. \square*

The moduli space of shrubs is topologised in exactly the same way as the Stasheff moduli space. In particular, it also admits a decomposition into smooth manifolds. One way to see this is to observe

Lemma B.5. *There is surjective projection*

$$\mathcal{S}_d \rightarrow \mathcal{T}_d$$

obtained by replacing all incoming edges with edges of infinite length. The fiber $\tilde{\mathcal{S}}_d([T])$ over each closed cell $\tilde{\mathcal{T}}_d([T])$ admits a unique smooth structure which makes this projection a fibration with fibre an open interval, and with trivialization given by the length of an external edge.

Proof. The only part that requires checking is that the different trivializations determine compatible smooth structures. This follows easily from the fact that the difference between such trivializations can be expressed in terms of the lengths of internal edges which are, indeed, smooth functions on the base. \square

It is now clear that \mathcal{S}_d is not compact, and that its non-compactness comes from two phenomena; internal edges can expand to infinite length as in \mathcal{T}_d , and external edges can collapse to have length zero.

Definition B.6. *The moduli space of **singular shrubs** $\overline{\mathcal{S}}_d$ is the moduli space of singular ribbon trees with d incoming edges which satisfy the following properties:*

- *The incoming edges have finite type, and*
- *the outgoing edge has infinite type, and*
- *all irreducible components are stable, and are either shrubs or are Stasheff trees (all external edges are infinite), and*

- the distance from any point to the incoming vertices lying above it is independent of the choice of incoming vertex.

Remark B.7. It is immediate from the above Definition that the irreducible components containing the incoming vertices are shrubs, and that all other irreducible components are Stasheff trees.

The main result of this Appendix is

Proposition B.8. $\overline{\mathcal{S}}_d$ is a compact manifold with having \mathcal{S}_d as its interior. Moreover, it admits a decomposition into smooth manifolds with corners $\overline{\mathcal{S}}_d[T]$ labeled by the topological type of the shrub.

We first observe

Lemma B.9. By replacing all incoming vertices with vertices of infinite type and collapsing any unstable components, we obtain a projection map

$$\mathbf{p} : \overline{\mathcal{S}}_d([T]) \rightarrow \overline{\mathcal{T}}_d([T]).$$

The fibre at a point $T \in \mathcal{T}_d$ is identified by the length function of an incoming edge e with a closed interval $[a_e, +\infty]$. \square

As we shall see presently, this lemma essentially determines the topology and smooth structure of each cell away from the inverse image of $\partial\overline{\mathcal{T}}_d$. We shall call the part of the boundary lying over \mathcal{T}_d the **horizontal boundary** and the part lying over the boundary of $\overline{\mathcal{T}}_d$ the **vertical boundary** of $\overline{\mathcal{S}}_d$. In order to understand this later subset, we consider

Definition B.10. A collection of successive vertices of a ribbon tree forms a **branch** if the complement of the minimal tree containing them is connected.

Definition B.11. An **order-preserving** partition of $\{1, \dots, d\}$ is a partition into subsets $A_1, \dots, A_{d'}$ such that $a_i < a_j$ if $a_i \in A_i$ and $a_j \in A_j$. We will denote the set of such partition into d' subsets by $\text{Part}(d, d')$.

Given a tree T' with d' incoming edges, and such a partition, we obtain a unique tree T with d incoming edges, such that those that belong to a subset that is not a singleton have vanishing length. We have

Lemma B.12. The horizontal boundary of $\overline{\mathcal{S}}_d([T])$ is covered by the images of the inclusions of

$$\overline{\mathcal{S}}_{d'}([T']) \times \text{Part}(d, d') \rightarrow \overline{\mathcal{S}}_d([T]),$$

and by the section of \mathbf{p} corresponding to all incoming edges having infinite length.

The vertical boundary of $\overline{\mathcal{S}}_d([T])$ is covered by the images of the inclusions for each $1 \leq n \leq d$ of

$$\coprod_{\sum_{1 \leq i \leq n} d_i = d} \overline{\mathcal{S}}_{d_1}([T_1]) \times \dots \times \overline{\mathcal{S}}_{d_n}([T_n]) \times \overline{\mathcal{T}}_n([T_0]) \rightarrow \overline{\mathcal{S}}_d([T]).$$

with the restriction that there exists some $d_i \neq 1$.

Proof. The section of \mathbf{p} takes a Stasheff tree with d inputs and attaches half infinite intervals (isometric to $[0, +\infty)$) to all incoming vertices. This covers the part of the horizontal boundary where some incoming edge (hence all incoming edges) reach infinite length. The remainder of the horizontal boundary corresponds to an edge collapsing to have length 0.

Given a tree T , we consider the largest subtree T_0 containing the outgoing vertex, but no binary vertex. If such a tree is included properly in T , then its complement is a disjoint union of trees T_i that occur in lower dimensional moduli spaces and T is in the image of the inclusion

$$\coprod_{\sum_{1 \leq i \leq n} d_i = d} \overline{\mathcal{S}}_{d_1}([T_1]) \times \cdots \times \overline{\mathcal{S}}_{d_n}([T_n]) \times \overline{\mathcal{T}}_n([T_0]) \rightarrow \overline{\mathcal{S}}_d([T]).$$

If $d_i = 1$ for all i , then each \mathcal{S}_{d_i} is a point and we are considering a component of the horizontal boundary. \square

One way of stating the naturality of the above boundary inclusions is to observe

Lemma B.13. *The following diagram commutes*

$$\begin{array}{ccc} \overline{\mathcal{S}}_{d_1}([T_1]) \times \cdots \times \overline{\mathcal{S}}_{d_n}([T_n]) \times \overline{\mathcal{T}}_n([T_0]) & \longrightarrow & \overline{\mathcal{S}}_d([T]) \\ \downarrow & & \downarrow \\ \overline{\mathcal{T}}_{d_1}([T_1]) \times \cdots \times \overline{\mathcal{T}}_{d_n}([T_n]) \times \overline{\mathcal{T}}_n([T_0]) & \longrightarrow & \overline{\mathcal{T}}_d([T]), \end{array}$$

where the vertical arrows are given by the various projections \mathbf{p} , and the bottom horizontal arrow is the usual gluing map which makes $\overline{\mathcal{T}}_d$ into a topological operad.

We need a few more pieces of notation. If (d_1, \dots, d_n) is a partition of d , let D_i denote the partial sums

$$D_i = \sum_{j \leq i} d_j.$$

For each tree T , let

$$T(D_i, D_{i+1})$$

denote the minimal connected subtree of T containing all vertices whose labels range from D_i to $D_{i+1} - 1$. Let $U(d_1, \dots, d_n)$ denote the subset of $\overline{\mathcal{S}}_d$ consisting of trees such that all $T(D_i, D_{i+1})$ are branches. Similarly, let $e(D_i, D_{i+1})$ denote the outgoing edge at the outgoing vertex of $T(D_i, D_{i+1})$ (In particular, $e(D_i, D_{i+1})$ is not contained in this tree).

We are now ready to give the

Proof of Proposition B.8. The smooth structure on $\overline{\mathcal{S}}_d[T]$ is already determined in the interior by Lemma B.5. Along the horizontal boundary, it can be defined by requiring that \mathbf{p} is a fibration, with trivialization given by the length function of an incoming edge along the component of the horizontal boundary covered by the images of

$$\overline{\mathcal{S}}_{d'}[T'] \times \text{Part}(d, d') \rightarrow \overline{\mathcal{S}}_d[T].$$

Along the other component of the horizontal boundary (where the length of incoming edges takes infinite length), we require that the smooth trivialization be given by the exponential of the negative length function of an incoming edge. Again, this smooth structure is independent of which incoming edge is chosen since the length of incoming edges differ from each other by the lengths of the internal edges, which are smooth functions on the base.

It suffices therefore to define the topology (and smooth structure) along the vertical boundary. We shall proceed by induction; in the base case, both $\overline{\mathcal{S}}_1$ and \mathcal{S}_1 are points. Our induction hypothesis is that

- Whenever the length of an edge is finite, it defines a smooth function to \mathbb{R} .
- When the length of an edge reaches infinity, then the exponential of its negative is a smooth function.
- The subsets of $\mathcal{S}_i([T])$ consisting of homeomorphic trees with no vertices of valency greater than 4 are open (such trees are not required to be combinatorially equivalent, i.e. adding binary vertices will not change the homeomorphism type).

This last condition might seem superfluous; its presence is a consequence of the fact that the fibre over a boundary point of $\overline{\mathcal{T}}_n$ can consist of rather complicated strata, and that we have not yet specified how these strata fit together.

By the previous lemma, every point T on the vertical boundary lies in the image of

$$\overline{\mathcal{S}}_{d_1}([T_1]) \times \cdots \times \overline{\mathcal{S}}_{d_n}([T_n]) \times \overline{\mathcal{T}}_n([T_0]) \rightarrow \overline{\mathcal{S}}_d([T]).$$

By induction we have a topology (and smooth structure as a manifold with corners) on the source of the above map. We shall require these maps to be diffeomorphisms onto their images. There is no obstruction to achieving this since the images of such maps intersect along products of lower dimensional moduli spaces whose smooth structure is already determined.

Consider the set $U(d_1, \dots, d_n)$ defined in the paragraph preceding the beginning of this proof. We impose the condition that these sets are open. This condition is, by induction, compatible with all previous choices. We define a map

$$\mathfrak{t} : U(d_1, \dots, d_n) \rightarrow \overline{\mathcal{S}}_{d_1}([T_1]) \times \cdots \times \overline{\mathcal{S}}_{d_n}([T_n]) \times \overline{\mathcal{T}}_n([T_0])$$

by replacing each edge $e(D_i, D_{i+1})$ by an edge of infinite length. This map is clearly surjective. Let $U(K, d_1, \dots, d_n)$ denote the open subset of $\overline{\mathcal{S}}_d([T])$ where all edges $e(D_i, D_{i+1})$ have length greater than K . If $V_0(\frac{1}{K})$ is a fundamental system of neighbourhoods of T_0 in the product

$$\overline{\mathcal{S}}_{d_1}([T_1]) \times \cdots \times \overline{\mathcal{S}}_{d_n}([T_n]) \times \overline{\mathcal{T}}_n([T_0]),$$

we define a fundamental system of neighbourhoods of T_0 in $\overline{\mathcal{S}}_d([T])$ to be

$$U_0\left(\frac{1}{K}\right) = \mathfrak{t}^{-1}\left(V_0\left(\frac{1}{K}\right)\right) \cap U(K, d_1, \dots, d_n).$$

Again, this is clearly compatible with previous choices of topology. To define the smooth structure as a manifold with boundary, we observe that the projection map

$$U_0\left(\frac{1}{K}\right) \rightarrow V_0\left(\frac{1}{K}\right)$$

is trivialized by the length function of an edge $e(D_i, D_{i+1})$ with target $(K, +\infty]$. We determine the smooth structure by again requiring that the exponential of the negative length is smooth. Since by induction, $V_0(\frac{1}{K})$ is diffeomorphic to the standard model of a neighbourhood of a corner, the same is true for its product with a closed interval. \square

APPENDIX C. ORIENTATIONS IN MORSE THEORY

In order to define an A_∞ structure on \mathbb{Z} -valued Morse complexes, we must orient the moduli spaces of gradient trees. The first step is to coherently orient the Stasheff polyhedra. This is, in fact, relatively easy to do using the well-known identification with the moduli space of discs. The reader can consult Section (12g) of [40] for

details. Note that such an orientation induces an orientation of each maximal cell $\mathcal{T}_d([T])$.

We now orient the moduli spaces of gradients trees which appear in the definition of the higher products in the Morse pre-category. Consider the manifold

$$\mathcal{E}(q, \vec{p}) = W^s(q) \times W^u(p_1, \dots, p_d) \times \mathcal{T}_d,$$

Let us assume that orientations for all descending manifolds have been chosen. If we write $\lambda(N)$ for the top exterior power of the tangent bundle of N as a graded vector space, then the induced orientation is given by the isomorphism

$$(C.1) \quad \lambda(\mathcal{T}(q, \vec{p})) \cong \lambda(\Delta^{d+1}) \otimes \lambda^{-1}(Q^{d+1}) \otimes \lambda(\mathcal{E}(q, \vec{p})),$$

where the tangent space $\mathcal{E}(q, \vec{p})$ at a given point is defined to be the tangent space of the corresponding cell $\tilde{\mathcal{E}}(q, \vec{p}; [T])$ for $[T]$ a trivalent tree; the global orientability of the Stasheff polyhedra implies that the choice of a resolution of a non-trivalent tree does not affect the signs.

We will twist this orientation by

$$(C.2) \quad \sigma(q, \vec{p}) = (m+1) \cdot \left(\sum_{j=0}^d (\deg(q) + \sum_{k \leq j} \deg(p_k)) + \dim(\mathcal{T}(q, \vec{p})) \cdot (1 + m + d + \deg(q)) \right).$$

Remark C.1. In Equation (C.1), we used the inverse of vector space $\lambda^{-1}(Q^{d+1})$. Our convention for orienting them is as follows: If $v_1 \wedge \dots \wedge v_n$ is an orientation of λV , then $v_n^{-1} \wedge \dots \wedge v_1^{-1}$ is the “inverse” orientation of $\lambda^{-1}V$.

We can now complete the

Proof of Proposition 4.22 . Let $\mathcal{T}(q, \vec{p})$ be a 1-dimensional moduli spaces of gradient trees. We know that its boundary consists of points corresponding to

$$\mathcal{T}(q, \vec{p}_1^i, r, \vec{p}_{i+d_2+1}^d) \times \mathcal{T}(r, \vec{p}_{i+1}^{i+d_2}),$$

where

$$\deg(r) = d_2 - 2 + \sum_{j=i+1}^{i+d_2} \deg(p_j),$$

and the sequences \vec{p}_a^b stand for $(p_a, \dots, p_{b-1}, p_b)$. Set $d_1 = d - d_2 + 1$. For simplicity, we assume that $d_2 \neq 1$, in other words that we’re working entirely with stable trees. The idea does not change in the case where $d_2 = 1$, though the notation is complicated by the presence of a factor associated to the group of automorphisms of moduli spaces of gradient trajectories between critical points.

It suffices to compare the induced orientation of the boundary with its orientation as a product. Ignoring the factors $\sigma(q, \vec{p})$ for now, we note that an orientation of a tubular neighbourhood of the boundary is given by

$$\begin{aligned} & \mathbb{R} \cdot \nu \otimes \lambda(\Delta^{d_1+1}) \otimes \lambda^{-1}(Q^{d_1+1}) \otimes \lambda(\mathcal{E}(q, \vec{p}_1^i, r, \vec{p}_{i+d_2+1}^d)) \\ & \otimes \lambda(\Delta^{d_2+1}) \otimes \lambda^{-1}(Q^{d_2+1}) \otimes \lambda(\mathcal{E}(r, \vec{p}_{i+1}^{i+d_2})), \end{aligned}$$

where ν is the outward normal vector at the boundary. One can pass from this orientation to the orientation as an open subset of $\mathcal{T}(q, \vec{p})$ through simple steps.

We go through them one by one and record the corresponding change of orientation. Note that all computations are done modulo 2.

- (1) Recall that the $d_1 + 1$ factors Q^{d_1+1} correspond to the output q and the inputs $(p_1, \dots, p_i, r, p_{i+d_2+1}, \dots, p_d)$. Writing $\lambda^{-1}(Q^{d_1+1})$ as the tensor power of the corresponding number of copies of $\lambda^{-1}(Q)$, we permute these lines so that the factor corresponding to r is adjacent to $\lambda(\Delta^{d_1+1})$. Because we're working with the inverse of top exterior power of the tangent bundle, we have to slide the factor corresponding to r past those corresponding to p_{i+d_2+1}, \dots, p_d . This yields a sign given by the parity of

$$(C.3) \quad m \cdot m(d - d_2 - i) = m(d - d_2 - i).$$

Since the diagonal is oriented in the same way as Q , we have a natural identification of

$$\lambda(\Delta^{d_1+1}) \otimes \lambda^{-1}(Q) \rightarrow \mathbb{R},$$

and we can cancel these factors.

- (2) Similarly, we permute the factors of Q^{d_2+1} so that the copy corresponding to r appears next to Δ^{d_2+1} , and cancel these two factors. The permutation introduces a sign

$$(C.4) \quad m \cdot md_2 = md_2.$$

- (3) We decompose

$$\mathcal{E}(q, \vec{p}_1^i, r, \vec{p}_{i+d_2+1}^d) = W^s(q) \times W^u(p_1, \dots, p_i) \times W^u(r) \times W^u(p_{i+d_2+1}, \dots, p_d) \times \mathcal{T}_{d_1},$$

and permute $W^u(r)$ with $Q^{d_1} \times W^s(q) \times W^u(p_1, \dots, p_i)$. This changes the orientation by

$$(C.5) \quad \begin{aligned} & (m - \deg(r)) \left(md_1 + \deg(q) + im + \sum_{j=1}^i \deg(p_j) \right) \\ &= (m - \deg(r)) (md_1 + \deg(q) + i(m+1) + \mathbf{\boxtimes}_i). \end{aligned}$$

- (4) We move the $W^s(r)$ factor of $\mathcal{E}(r, \vec{p}_{i+1}^{i+d_2})$ past

$$Q^{d_1} \times W^s(q) \times W^u(p_1, \dots, p_i, p_{i+d_2+1}, \dots, p_d) \times \mathcal{T}_{d_1} \times Q^{d_2}$$

with a sign of

$$(C.6) \quad \begin{aligned} & \deg(r) \left(d_1 + \deg(q) + m - \sum_{j=1}^i \deg(p_j) - \sum_{j=i+d_2+1}^d \deg(p_j) + md_2 \right) \\ &= \deg(r) (d_1 + \deg(q) + m + d_1 + \deg(q) + \deg(r) + md_2) \\ &= \deg(r) (md_2 + m + 1). \end{aligned}$$

- (5) We now permute Q^{d_2} with the terms that separate it from Q^{d_1} . The corresponding sign is

$$(C.7) \quad \begin{aligned} & md_2 \left(d_1 + \deg(q) + m(d_1 - 1) - \sum_{j=1}^i \deg(p_j) + \sum_{j=i+d_2+1}^d \deg(p_j) \right) \\ &= md_2 (m(d_1 - 1) + \deg(r)). \end{aligned}$$

- (6) Recall that the factors of Q^{d+1} correspond to the sequence (q, p_1, \dots, p_d) . The terms of $Q^{d_1} \times Q^{d_2}$ are not in this same order. Since we're taking the inverse of the canonical bundle, the desired permutation has sign

$$(C.8) \quad md_2(i+1).$$

Let us note that, after taking into account the signs which we introduced in the previous steps, our orientation of a neighbourhood of the boundary is now given by

$$\mathbb{R} \cdot \nu \otimes \lambda(W^u(r)) \otimes \lambda(W^s(r)) \otimes \lambda^{-1}(Q^{d+1}) \otimes \lambda(W^s(q)) \otimes \lambda(W^u(p_1, \dots, p_i, p_{i+d_2+1}, \dots, p_d)) \otimes \lambda(\mathcal{T}_{d_1}) \otimes \lambda(W^u(p_i+1, \dots, p_{i+d_2})) \otimes \lambda(\mathcal{T}_{d_2}).$$

We chose our orientation of $W^u(r)$ in such a way as to have an oriented isomorphism

$$\lambda(W^u(r)) \otimes \lambda(W^s(r)) \cong \lambda(Q).$$

The next steps to arrive at our chosen orientation of $\mathcal{T}(q, \vec{p})$ are:

- (7) Permute $W^u(p_i+1, \dots, p_{i+d_2})$ with $W^u(p_{i+d_2+1}, \dots, p_d) \times \mathcal{T}_{d_1}$, and introduce a sign

$$(C.9) \quad \left(d_2 m - \sum_{j=i+1}^{i+d_2} \deg(p_j) \right) \cdot \left(d_1 + (d - d_2 - i)m + \sum_{j=i+d_2+1}^d \deg(p_j) \right) \\ = (d_2(m+1) + \deg(r)) (i + (d - d_2 - i)m + \deg(q) + \deg(r) + \mathfrak{X}_i)$$

- (8) Permute the normal vector $\mathbb{R} \cdot \nu$ to be adjacent to $\mathcal{T}_{d_1} \times \mathcal{T}_{d_2}$. The corresponding sign is given by

$$(C.10) \quad dm + \deg(q) + dm + \sum_{j=1}^d \deg p_j = d+1$$

The total sign change so far is given by the parity of

$$(\deg(q) + \mathfrak{X}_i)(m + d_2(m+1)) \\ + \deg(r)(m+1)d_2 + md_2d + mi + md_2 + d_2i + d + m + 1$$

There are two more sources of signs that we have neglected.

- (9) There is a choice of coherent orientations on the Stasheff polyhedra such that the orientation of a neighbourhood of the boundary of \mathcal{T}_d differs from the product orientation

$$\mathbb{R} \cdot \nu \times \mathcal{T}_{d_1} \times \mathcal{T}_{d_2}$$

precisely by

$$d_2(d-i) + i + d_2,$$

as explained in Equation (12.25) of [40].

- (10) The additional change of orientation σ given by Equation (C.2) introduces a sign equal to

$$(m+1)((d_2+1)(\deg(q) + \mathfrak{X}_i) + d_2 \deg(r) + dd_2 + d_2 + i) + 1 + m + d + \deg(q)$$

as can be easily checked by comparing the contribution of the inputs and outputs to σ in each of the three moduli spaces we're considering.

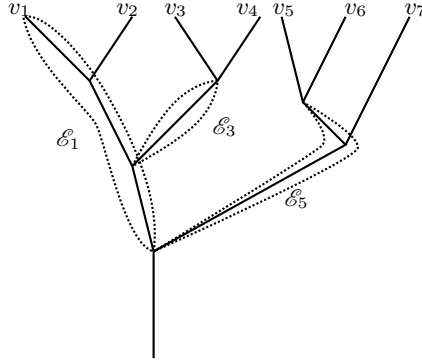


FIGURE 9.

The reader can easily check that the total sign change is, indeed, \mathfrak{X}_i . Since the (signed) sum of the number of points on the boundary of a 1-dimensional manifold vanishes, we conclude that the A_∞ equation holds. \square

We note that the complete proof of Proposition 4.32 also require orientations of moduli spaces of gradient trees, this time

$$\mathcal{S}(p, \vec{C}),$$

for a critical point p and cycles $\vec{C} = (C_1, \dots, C_d)$. In fact, such an orientation can be obtained through means entirely analogous to those used in the previous proof once we explain how to orient the moduli spaces of incoming-finite trees \mathcal{S}_d . We note that that the method we're about to use also gives orientations of the Stasheff polyhedra.

Consider a tree $T \in \mathcal{S}_d$. Let $\{v_i\}$ denote the incoming vertices of T . Let A_i be the descending arc containing v_i which is maximal with respect to the property of containing only finite edges. Let \mathcal{E}_1 be the ordered collection of finite edges whose union is A_1 , with the inverse order given by the natural orientation of the tree, i.e. the last element of \mathcal{E}_1 is the incoming edge at v_1 . Inductively define \mathcal{E}_i as the ordered collection of edges of A_i , excluding the incoming edges if $i \geq 2$, with the additional restriction that the edges appearing in \mathcal{E}_i do not appear in any preceding set \mathcal{E}_j . We write $|\mathcal{E}_i|$ for the number of elements in such a collection. See Figure 9.

Lemma C.2. *If T is trivalent, the lengths of the edges which appear in*

$$\bigcup_i \mathcal{E}_i$$

give local coordinates for \mathcal{S}_d near T . \square

Note that some of the collections \mathcal{E}_i may be empty. We now use this canonical choice of coordinates to define an orientation. First, we need

Definition C.3. *If e is an edge of a trivalent tree T , then the **dexterity** of e , denoted $r_T(e)$ is the number of right turns along the geodesic from e to the outgoing vertex.*

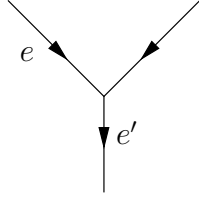


FIGURE 10. The path from e to the node of T takes a right turn at e'

In other words, if A_e is the geodesic from e to the outgoing vertex, then we can stand at every edge of A_e , looking towards the outgoing vertex, and decide whether the next edge of A_e will correspond to a left turn or a right turn at the next vertex (See Figure 10). Left and right are, of course, determined by the chosen cyclic ordering at each vertex.

To simplify notation, we write de for $dl(e)$. We define an orientation by

$$\Omega_T(v_i) = \bigwedge_{e \in \mathcal{E}_i} (-1)^{r_T(e)} de$$

$$\Omega_T = \Omega_T(v_1) \wedge \cdots \wedge \Omega_T(v_d)$$

where the edges that appear in the first expression are ordered along the path from the outgoing vertex to v_i .

Lemma C.4. *The orientation Ω_T is independent of the topological type of the tree T .*

Proof. It suffices to consider neighbourhoods of trees with one quadrivalent vertex (these form hyperplanes in the moduli space of trees separating the open cells consisting of trivalent trees of a fixed topological type). Crossing such a hyperplane, a branch $B(j, k)$ of the tree moves from one side of a trivalent vertex to the other. Such a branch is uniquely determined by its incoming leaves $\{v_j, \dots, v_k\}$. See Figure 11.

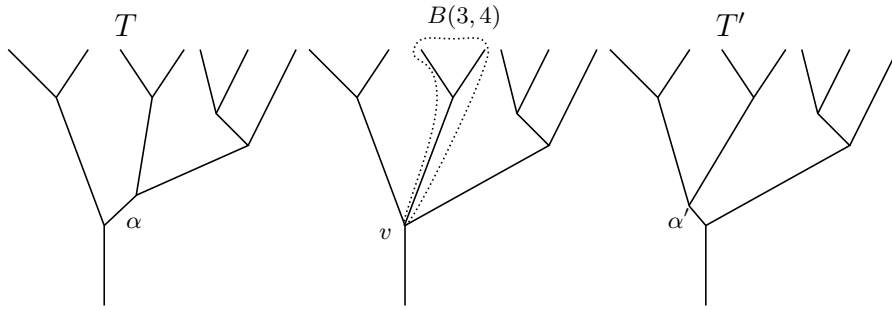


FIGURE 11.

Fix a tree with a unique quadrivalent vertex v . Note that all our coordinates extend naturally through the hyperplane, except for the map to \mathbb{R} which is given by the length of the edge which is collapsed at v . Let T and T' be nearby trees on either side of the hyperplane. We fix a co-orientation of the hyperplane by assuming that the path from v_{k+1} to the outgoing vertex includes the edge $\alpha \subset T$

which is collapsed when passing through the hyperplane. While the corresponding edge $\alpha' \subset T'$ leads, by necessity, to v_{j-1} . In fact, we can extend $l(\alpha)$ smoothly to the other side of the hyperplane by declaring it equal to $-l(\alpha')$. In particular, we have

$$\begin{aligned} d\alpha &= -d\alpha' \\ r_T(\alpha) &= 1 + r_{T'}(\alpha'). \end{aligned}$$

Let $B(i, \ell)$ denote the branch of T containing v as its bottom-most vertex. It is obvious that for every edge that does not lie on $B(i, \ell)$, we have

$$r_T(e) = r_{T'}(e).$$

We claim the same equality holds for all edge on $B(k+1, \ell)$ and for $B(j, k)$. Indeed, we only remove a left turn from the paths with origin the edges of $B(k+1, \ell)$ to the outgoing vertex, while we permute the order of a left and right turn on corresponding paths from the edges of $B(j, k)$. It suffices therefore to check that

$$\Omega_T(v_i) \wedge \cdots \wedge \Omega_T(v_j) = \Omega_{T'}(v_i) \wedge \cdots \wedge \Omega_{T'}(v_j).$$

Observe $d\alpha$ is the first 1-form which appears in our formula for $\Omega_T(v_j)$, while $d\alpha'$ appears in our formula for $\Omega_{T'}(v_i)$. Let us therefore split

$$\Omega_{T'}(v_i) = \Omega_{T'}^1(v_i) \wedge d\alpha' \wedge \Omega_{T'}^2(v_i).$$

The result now follows easily by observing that the dexterity of each edge in $B(i, j-1)$, changes by 1 as α is collapsed, which exactly cancels the sign produced by commuting $d\alpha$ past

$$\Omega_T^2(v_i) \wedge \cdots \wedge \Omega_T(v_{j-1}).$$

□

APPENDIX D. A_∞ PRE-CATEGORIES

We recall some elementary definitions concerning A_∞ pre-categories, due to Kontsevich and Soibelman [28].

Definition D.1. *An A_∞ pre-category \mathcal{A} consists of the following data:*

- A collection of objects $\mathcal{Ob}(\mathcal{A})$.
- For each integer n , a collection of transverse sequences

$$\mathcal{Ob}_{tr}^n(\mathcal{A}) \subset \mathcal{Ob}(\mathcal{A})^n$$

- For each ordered pair $(X_0, X_1) \in \mathcal{Ob}_{tr}^2(\mathcal{A})$, a \mathbb{Z} -graded chain complex

$$\mathcal{A}(X_0, X_1).$$

- For each transverse sequence of length $d+1$, (X_0, X_1, \dots, X_d) , a degree $d-2$ map of graded vector spaces, called the d th higher product

$$\mathcal{A}(X_{d-1}, X_d) \otimes \cdots \otimes \mathcal{A}(X_0, X_1) \xrightarrow{\mu_d} \mathcal{A}(X_0, X_d).$$

The defining data of an A_∞ pre-category is required to satisfy the following compatibility conditions

- Every subsequence of a transverse sequence is transverse.
- The higher products are required to satisfy the A_∞ equation, i.e. Equation (3.3).

Kontsevich and Soibelman further defined the notion of quasi-isomorphisms in A_∞ pre-categories, and used it to give a criterion of unitality of A_∞ pre-categories.

Definition D.2. A morphism $f \in \mathcal{A}(X_0, X_1)$ is said to be a **quasi-isomorphism** if, whenever (Z, X_0, X_1) , or (X_0, X_1, Y) are transverse triples, the maps

$$\begin{aligned} \mathcal{A}(X_1, Y) &\xrightarrow{\mu_2(-, f)} \mathcal{A}(X_0, Y) \\ \mathcal{A}(Z, X_0) &\xrightarrow{\mu_2(f, -)} \mathcal{A}(Z, X_1) \end{aligned}$$

induce isomorphisms on homology.

Definition D.3. An A_∞ pre-category is unital if for each object X , and each transverse sequence S , there exist objects X_+ and X_- which are quasi-isomorphic to X such that the sequence

$$(X_-, S, X_+)$$

is transverse.

In this setting, it is appropriate to replace the usual notion of full subcategory.

Definition D.4. The inclusion of a full sub-pre-category \mathcal{B} in an A_∞ pre-category \mathcal{A} is said to be **essentially surjective** if every object of \mathcal{A} is quasi-isomorphic to an object of \mathcal{B} .

In order to simplify our notation, we deviate slightly from Kontsevich and Soibelman's definition of A_∞ functor.

Definition D.5. An A_∞ functor F between A_∞ pre-categories \mathcal{A} and \mathcal{B} consists of the following data:

- An essentially surjective inclusion $\mathcal{A}_F \subset \mathcal{A}$.
- A map

$$F : \mathcal{O}b(\mathcal{A}_F) \rightarrow \mathcal{O}b(\mathcal{B}).$$

- For each transverse sequence (X_0, \dots, X_d) of \mathcal{A}_F , a degree $d-1$ map

$$F^d : \mathcal{A}(X_{d-1}, X_d) \otimes \dots \otimes \mathcal{A}(X_0, X_1) \rightarrow \mathcal{B}(FX_0, FX_d).$$

This data must satisfy the following conditions

- The image of every transverse sequence of \mathcal{A}_F is transverse in \mathcal{B} .
- The maps F^d satisfy the A_∞ equation for functors

$$\begin{aligned} \sum_{l_1 + \dots + l_k = d} \mu_k^B \circ (F^{l_1} \otimes \dots \otimes F^{l_k}) = \\ \sum_{i+d_2 \leq d} (-1)^{\mathfrak{F}_i} F^{d-d_2+1} \circ (\mathbb{1}^{d-i-d_2} \otimes \mu_{d_2}^A \otimes \mathbb{1}^i). \end{aligned}$$

This definition differs from that of [28] where they require $\mathcal{A}_F = \mathcal{A}$. However, this is irrelevant since this difference disappears upon considering quasi-equivalence classes of A_∞ pre-categories.

Definition D.6. An A_∞ **quasi-equivalence** is an A_∞ functor which induces an isomorphism at the level of homology categories.

Two A_∞ pre-categories A and B are **equivalent** if there exists a sequence $\{A_i\}_{i=0}^n$ of A_∞ pre-categories and A_∞ quasi-equivalences

$$A = A_0 \rightarrow A_1 \leftarrow A_2 \rightarrow A_3 \leftarrow \dots \rightarrow A_n = B.$$

APPENDIX E. CELLULAR CHAINS AND SIMPLICIAL COCHAINS ON MANIFOLDS
WITH BOUNDARY

Let Q be a smooth compact oriented manifold with boundary ∂Q . We would like to rephrase the classical equivalence between the cup product on the simplicial cochains of a triangulation \mathcal{P} and the intersection pairing on cellular chains in a language more suitable to our goals in this paper. A stronger result appears in [31] (the commutative structure is taken into account, cf. [44]), but the techniques used seem incompatible with our specific setup. To make the definition of the cup product standard, we work with a special type of triangulation:

Definition E.1. A triangulation \mathcal{P} is **simplicial** if there is a complete ordering on the vertices of \mathcal{P} such that every cell of \mathcal{P} can be expressed uniquely as the “hull” of an ordered sequence of vertices $v_0 < \dots < v_k$. We write $[v_0, \dots, v_k] \subset \mathcal{P}$ for such a cell.

With this condition, one can use the full machinery of simplicial sets, though we shall not. We now construct a category which encodes the cup product on relative simplicial cochains.

Definition E.2. The pre-category $\text{Simp}(\mathcal{P})$ has as objects functions

$$H: \partial Q \rightarrow \mathbb{R}.$$

A pair H_0 and H_1 of such functions is transverse if the difference $H = H_1 - H_0$ satisfies the following condition: the inclusion of $H^{-1}([0, +\infty))$ into $\partial_H^+ \mathcal{P}$, the union of all cells where H takes some positive value, is a deformation retract.

If H_0 and H_1 are transverse, the space of morphisms from H_0 to H_1 is given by

$$\text{Simp}(\mathcal{P})(H_0, H_1) = C^*(\mathcal{P}, \partial_{H_1 - H_0}^+ \mathcal{P}).$$

A sequence (H_0, H_1, \dots, H_d) is transverse if all pairs (H_i, H_j) with $i < j$ are transverse. All higher products vanish, and the composition is defined through the inclusion $\partial_{H_2 - H_0}^+ \mathcal{P} \subset \partial_{H_2 - H_1}^+ \mathcal{P} \cup \partial_{H_1 - H_0}^+ \mathcal{P}$ and the corresponding diagram:

$$\begin{array}{ccc} C^*(\mathcal{P}, \partial_{H_2 - H_1}^+ \mathcal{P}) \otimes C^*(\mathcal{P}, \partial_{H_1 - H_0}^+ \mathcal{P}) & \longrightarrow & C^*(\mathcal{P}, \partial_{H_2 - H_1}^+ \mathcal{P} \cup \partial_{H_1 - H_0}^+ \mathcal{P}) \\ & & \downarrow \\ & & C^*(\mathcal{P}, \partial_{H_2 - H_1}^+ \mathcal{P}). \end{array}$$

For the sake of definiteness, we describe the chain model theory that we use. Our treatment is standard in the sense that it can be extracted from an algebraic topology textbook [23, 30]. The chain complex of a simplicial triangulation is defined by

$$C_k(\mathcal{P}) \equiv \bigoplus_{[v_0, \dots, v_k]} \mathbb{Z} \cdot [v_0, \dots, v_k].$$

Again, we reiterate the fact that our conventions are such that writing $[v_0, \dots, v_k]$ subsumes the fact that the vertices are listed in increasing order. The differential is given as usual by

$$\partial[v_0, v_1, \dots, v_k] = \sum_{i=0}^k (-1)^i [v_0, \dots, v_{i-1}, \hat{v}_i, v_{i+1}, \dots, v_k].$$

If $A \subset \mathcal{P}$ is a subcomplex, we define the relative cochain complex

$$C^*(\mathcal{P}, A) \equiv \{\phi: C_*(\mathcal{P}) \rightarrow \mathbb{Z} \mid \phi(\sigma) = 0 \text{ if } \sigma \in A\}.$$

For subcomplexes $A, B \subset \mathcal{P}$, we define the cup product

$$(E.1) \quad C^*(\mathcal{P}, A) \otimes C^*(\mathcal{P}, B) \rightarrow C^*(\mathcal{P}, A \cup B)$$

by the formula

$$(E.2) \quad \phi \cup \psi([v_0, \dots, v_k]) = \sum_{i=0}^k \phi([v_0, \dots, v_i]) \cdot \psi([v_i, \dots, v_k]).$$

A priori, $\phi \cup \psi$ is a homomorphism from $C_*(\mathcal{P})$ to \mathbb{Z} , but, since every cell of $A \cup B$ lies either in A or in B , it is easy to check that $\phi \cup \psi$ vanishes on all such cells and hence defines an element of $C^*(\mathcal{P}, A \cup B)$ (cf. [23, p. 209] where the less straightforward case of singular cohomology is discussed).

We now proceed with a cellular version of this construction. By a cellular subdivision \mathcal{R} of Q , we mean a covering by n -dimensional closed polytopes such that the intersection of any two polytopes is either empty or a common cell on the boundary. Given such a subdivision, we define the corresponding cellular chain complex by

$$C_*(\mathcal{R}) \equiv \bigoplus_{C \in \mathcal{R}} \mathbb{Z} \cdot \Omega(C),$$

where $\Omega(C)$ is the orientation line of the cell C . The differential is obtained by restricting the orientation to the boundary. Let us fix a subdivision $\check{\mathcal{P}}$ dual to a simplicial triangulation \mathcal{P} , and a corresponding constant $\epsilon > 0$ which is much smaller than the diameter of all cells in $\check{\mathcal{P}}$. It will be convenient to use the following terminology:

Definition E.3. *The ϵ -essential part $C - \partial C(\epsilon)$ of a cell $C \in \check{\mathcal{P}}$ is the complement, in C , of an ϵ -neighbourhood of ∂C .*

By definition, every cell C of codimension k in $\check{\mathcal{P}}$ is dual to a cell of the form $[v_0, \dots, v_k]$, and the link of C is isomorphic to the simplex with vertices (v_0, \dots, v_k) . We will sometimes write σ for a cell of \mathcal{P} and $\check{\sigma}$ for the dual cell. Concretely, a point in the interior of C lies on the closure of exactly k top dimensional cells which are dual to the vertices $\{v_0, \dots, v_k\}$. We will have to understand cell subdivisions which are C^1 close to $\check{\mathcal{P}}$. To do this, it is helpful to first understand the local model.

Let Π^d be the cellular subdivision of \mathbb{R}^d into the $d + 1$ cells, d of which are of the form

$$\{(x_1, \dots, x_d) \mid 0, x_1, \dots, x_d \leq x_j\}$$

for some j between 1 and d , and the remaining cell is the set

$$\{(x_1, \dots, x_d) \mid x_1, \dots, x_d \leq 0\}.$$

It is easy to check that Π^d is a model for the dual subdivision of a d -simplex $\Delta^d = [v_0, \dots, v_d]$ with the cell where 0 dominates dual to v_0 , and the cell where x_j dominates all the other terms dual to the vertex v_j for $1 \leq j \leq d$. Given any positive real number r , we may construct a new cellular subdivision of \mathbb{R}^d by translating Π_d in the direction of the vector

$$(E.3) \quad (-r/d, \dots, -r/2, -r).$$

We write $\Pi^d(r)$ for this new subdivision of \mathbb{R}^d . See Figure 12. If σ is a cell of Δ^d , we write $\check{\sigma}$ for the dual cell of Π^d and $\check{\sigma}(r)$ for the dual cell of $\Pi^d(r)$.

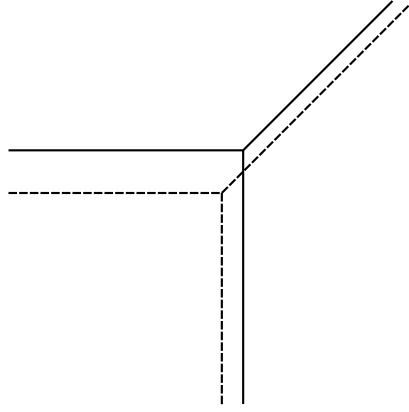


FIGURE 12. The one-skeleta of the polyhedral subdivisions Π^2 (full lines) and $\Pi^2(r)$ (dashed)

Lemma E.4. *If $\sigma = [u_0, \dots, u_k]$ and $\tau = [w_0, \dots, w_l]$ are cells of Δ^d then their duals $\check{\sigma}$ and $\check{\tau}(r)$ have non-empty intersection if and only if $u_k \leq w_0$.*

Proof. Observe that $\check{\sigma}$ and $\check{\tau}(r)$ intersect if and only if $[\check{u}_i]$ intersects $[\check{w}_j](r)$ for each i and j . However, it is easy to check that this can only happen if $u_i \leq w_j$, for all i and j , which because of our ordering conventions, is equivalent to $u_k \leq w_0$. \square

For any $r < r'$, the composition of translation by the vector of Equation (E.3) and a dilation of each variable maps $\Pi^d(r)$ to Π and $\Pi^d(r)$ to $\Pi^d(1)$. More generally, given an increasing collection r_i , the subdivisions $\Pi^d(r_i)$ are mutually transverse. Moreover, there is a self diffeomorphism of \mathbb{R}^d mapping $\Pi^d(i)$ to $\Pi^d(r_i)$, which can be easily made to depend smoothly on the parameters r_i . We conclude

Lemma E.5. *If σ^i is a sequence of cells of Δ^d , there exists a map*

$$(E.4) \quad \bigcap_i \check{\sigma}^i(i) \times \{(r_1, \dots, r_k) | r_1 \leq \dots \leq r_k\} \rightarrow \mathbb{R}^d$$

whose restriction to a fixed (r_1, \dots, r_k) is a diffeomorphism onto

$$\bigcap_i \check{\sigma}^i(r_i).$$

\square

To interpret the above results algebraically, we observe the intersections of all open cells in Π^d and $\Pi^d(r)$ give a new polyhedral subdivision of \mathbb{R}^d which we write $\Pi^d \pitchfork \Pi^d(r)$. Lemma E.4 implies that $\Pi^d \pitchfork \Pi^d(r)$ admits a canonical cellular map to Π^d which collapses

$$[\check{v}_i] \cap [\check{v}_j](r)$$

to $[v_i, v_j]$. It is easy to check that this map, which we've defined on the top simplices of $\Pi^d \pitchfork \Pi^d(r)$, extends to a cellular map. We write coll_Π for this map.

By composing the corresponding map on cellular chains with the intersection product

$$C_{n-i}(\Pi^d) \otimes C_{n-j}(\Pi^d(r)) \rightarrow C_{n-i-j}(\Pi^d \pitchfork \Pi^d(r))$$

we obtain an intersection product valued in $C_{n-i-j}(\Pi^d(r))$.

Lemma E.6. *The following diagram, where the vertical arrows are the canonical isomorphisms, is commutative*

$$\begin{array}{ccc} C_{n-i}(\Pi^d) \otimes C_{n-j}(\Pi^d(r)) & \longrightarrow & C_{n-i-j}(\Pi^d(r)) \\ \downarrow & & \downarrow \\ C^i(\Delta^d) \otimes C^j(\Delta^d) & \xrightarrow{\cup} & C^{i+j}(\Delta^d) \end{array}$$

Proof. This is essentially the content of Lemma E.4, once one observes that all intersections between cells of Π^d and $\Pi^d(r)$ that are not recorded by the cup products on $C^*(\Delta)$ are collapsed by the map $\Pi^d \cap \Pi^d(r) \rightarrow \Pi^d(r)$ to lower dimensional cells. \square

This local model readily generalizes to our original setting with the simplex $\Delta(d)$ replaced by a triangulation \mathcal{P} of a manifold Q .

Definition E.7. *A pair of subdivisions $(\check{\mathcal{P}}_0, \check{\mathcal{P}}_1)$ which are ϵ close to $\check{\mathcal{P}}$ in the C^1 topology are said to be **in simplicial position** if, whenever $C \in \check{\mathcal{P}}$ is a cell of dimension d , there exists a coordinate chart near the ϵ -essential part of C*

$$(C - \partial C(\epsilon)) \times \mathbb{R}^d \rightarrow Q$$

taking $(C - \partial C(\epsilon)) \times \Pi_0^d$ to $\check{\mathcal{P}}$, and $(C - \partial C(\epsilon)) \times \Pi^d(r)$ to $\check{\mathcal{P}}_1$, and which is compatible with the respective stratifications.

It is easy to check that constructing subdivisions in simplicial position is a completely local problem. In fact, starting with a subdivision $\check{\mathcal{P}}$, we use the ordering of the vertices of \mathcal{P} to order the top strata of $\check{\mathcal{P}}$. We may then construct a vector field which, on every cell of $\check{\mathcal{P}}$ points in the direction of the nearby top cells which appears first in this order. The image of $\check{\mathcal{P}}$ under the flow of this vector field (for any sufficiently small time) is in simplicial position with respect to $\check{\mathcal{P}}$.

If $(\check{\mathcal{P}}_0, \check{\mathcal{P}}_1)$ are in simplicial position, then the local copies of the polyhedral subdivisions $\Pi^d \cap \Pi^d(r)$ patch together globally to a cellular decomposition $\check{\mathcal{P}}_0 \cap \check{\mathcal{P}}_1$ of Q . The proof of the following global result is an immediate consequence of the local properties of Π^d and $\Pi^d(r)$ recorded in Lemmas E.4 and E.6:

Proposition E.8. *If $(\check{\mathcal{P}}_0, \check{\mathcal{P}}_1)$ are in simplicial position, then cells $\check{\sigma}_0 \in \check{\mathcal{P}}_0$ of and $\check{\tau}_1 \in \check{\mathcal{P}}_1$ dual respectively to $\sigma = [u_0, \dots, u_k]$ and $\tau = [w_0, \dots, w_l]$ intersect if and only if*

- $u_k = w_0$ and the cell $\rho = [u_0, \dots, u_k = w_0, \dots, w_l]$ exists, or
- $u_i < w_j$ for all i, j and the cell $\rho = [u_0, \dots, u_k, w_0, \dots, w_l]$ exists.

Moreover, if $\check{\mathcal{P}}_0$ and $\check{\mathcal{P}}_1$ are δ -close in the C^1 topology, then, in the first case, $\check{\sigma}_0 \cap \check{\tau}_1$ is δ -close to $\check{\rho}_1$ in the C^1 -topology (in particular, the cells are diffeomorphic), while $\check{\sigma}_0 \cap \check{\tau}_1$ lies in a δ -neighbourhood of the lower dimensional cell $\check{\rho}_1$ in the second case.

Finally, collapsing all cells of $\check{\mathcal{P}}_0 \cap \check{\mathcal{P}}_1$ of the second type defines a cellular map

$$\text{coll}_{\mathcal{P}_1}: \check{\mathcal{P}}_0 \cap \check{\mathcal{P}}_1 \rightarrow \check{\mathcal{P}}_1.$$

\square

As in the local model, the intersection pairing therefore defines a map

$$C_{n-i}(\check{\mathcal{P}}_0) \otimes C_{n-j}(\check{\mathcal{P}}_1) \rightarrow C_{n-i-j}(\check{\mathcal{P}}_0 \cap \check{\mathcal{P}}_1).$$

By composing the above map with the collapsing map $\check{\mathcal{P}}_0 \pitchfork \check{\mathcal{P}}_1 \rightarrow \check{\mathcal{P}}_1$, we obtain a product valued in $C_{n-i-j}(\check{\mathcal{P}}_1)$. We now have all the necessary pieces to introduce a cellular analogue of the category $\text{Simp}(\mathcal{P})$.

Definition E.9. *The category $\text{Cell}(\check{\mathcal{P}})$ has as objects pairs $(\check{\mathcal{P}}_0, H_0)$ where $\check{\mathcal{P}}_0$ is a cellular subdivision C^1 -close to $\check{\mathcal{P}}$, and H_0 is real valued function on ∂Q .*

A pair of objects $(\check{\mathcal{P}}_0, H_0)$ and $(\check{\mathcal{P}}_1, H_1)$ is said to be transverse if $(\check{\mathcal{P}}_0, \check{\mathcal{P}}_1)$ are in simplicial position and (H_0, H_1) is a transverse pair in $\text{Simp}(\mathcal{P})$. In this case, the space of morphisms between them is defined to be

$$\text{Cell}(\check{\mathcal{P}})_*((\check{\mathcal{P}}_0, H_0), (\check{\mathcal{P}}_1, H_1)) \equiv C_{n-*}(\check{\mathcal{P}}_1 - \partial_{H_1-H_0}^+ \check{\mathcal{P}}_1)$$

where

$$\partial_{H_1-H_0}^+ \check{\mathcal{P}}_1 = \{C \in \check{\mathcal{P}}_1 \mid C \cap \sigma \neq \emptyset, \text{ for some } \sigma \in \partial_{H_1-H_0}^+ \mathcal{P}\}.$$

A sequence of objects $\{(\check{\mathcal{P}}_i, H_i)\}_{i=0}^d$ is transverse whenever all pairs $(\check{\mathcal{P}}_i, \check{\mathcal{P}}_j)$ with $i < j$ are in simplicial position and the sequence $\{H_i\}_{i=0}^d$ is transverse in $\text{Simp}(\mathcal{P})$. All higher products vanish, and the multiplication for a transverse triple $(\check{\mathcal{P}}_0, H_0)$, $(\check{\mathcal{P}}_1, H_1)$, and $(\check{\mathcal{P}}_2, H_2)$

$$C_{n-j}(\check{\mathcal{P}}_2 - \partial_{H_2-H_1}^+ \check{\mathcal{P}}_2) \otimes C_{n-i}(\check{\mathcal{P}}_1 - \partial_{H_1-H_0}^+ \check{\mathcal{P}}_1) \rightarrow C_{n-i-j}(\check{\mathcal{P}}_2 - \partial_{H_2-H_0}^+ \check{\mathcal{P}}_2),$$

is defined by collapsing the intersection of cells in $\check{\mathcal{P}}_1$ and $\check{\mathcal{P}}_2$ to $\check{\mathcal{P}}_2$.

Proposition E.10. *The forgetful map $(\check{\mathcal{P}}_0, H_0) \rightarrow H_0$ extends to an A_∞ functor*

$$\text{Cell}(\check{\mathcal{P}}) \rightarrow \text{Simp}(\mathcal{P})$$

which is a quasi-equivalence.

Proof. At the level of morphisms, the functor simply identifies a cellular chain with the corresponding dual cochain. By construction, this is always an isomorphism of chain complexes. It suffices therefore to check that the functor respects composition; this is the content of Lemma E.6 locally, and can be readily seen to hold globally. \square

REFERENCES

- [1] Mohammed Abouzaid, *Homogeneous coordinate rings and mirror symmetry for toric varieties*, Geom. Topol. **10** (2006), 1097–1157 (electronic). MR **2240909** (2007h:14052)
- [2] Denis Auroux, Ludmil Katzarkov, and Dmitri Orlov, *Mirror symmetry for weighted projective planes and their noncommutative deformations*, Ann. of Math. (2) **167** (2008), no. 3, 867–943. MR 2415388
- [3] ———, *Mirror symmetry for del Pezzo surfaces: vanishing cycles and coherent sheaves*, Invent. Math. **166** (2006), no. 3, 537–582. MR **2257391** (2007g:14045)
- [4] Marston Morse and Emilio Baiada, *Homotopy and homology related to the Schoenflies problem*, Ann. of Math. (2) **58** (1953), 142–165. MR 0056922 (15,146g)
- [5] Martin Betz and Ralph L. Cohen, *Graph moduli spaces and cohomology operations*, Turkish J. Math. **18** (1994), no. 1, 23–41. MR **1270436** (95i:58037)
- [6] Alexei Bondal and Wei-Dong Ruan, *Mirror symmetry for weighted projective spaces*, in preparation.
- [7] Philip Candelas, Xenia C. de la Ossa, Paul S. Green, and Linda Parkes, *A pair of Calabi-Yau manifolds as an exactly soluble superconformal theory*, Nuclear Phys. B **359** (1991), no. 1, 21–74. MR **1115626** (93b:32029)
- [8] Vladimir Drinfeld, *DG quotients of DG categories*, J. Algebra **272** (2004), no. 2, 643–691. MR **2028075** (2006e:18018)
- [9] Tobias Ekholm, *Morse flow trees and Legendrian contact homology in 1-jet spaces*, Geom. Topol. **11** (2007), 1083–1224. MR **2326943** (2008i:53128)

- [10] Yakov Eliashberg and Mikhael Gromov, *Convex symplectic manifolds*, Several complex variables and complex geometry, Part 2 (Santa Cruz, CA, 1989), 1991, pp. 135–162. MR1128541 (93f:58073)
- [11] Bohan Fang, Chiu-Chu Melissa Liu, David Treumann, and Eric Zaslow, *T-Duality and Equivariant Homological Mirror Symmetry for Toric Varieties*, available at [arXiv:0811.1228](#).
- [12] Andreas Floer, *Morse theory for Lagrangian intersections*, J. Differential Geom. **28** (1988), no. 3, 513–547. MR965228 (90f:58058)
- [13] Andreas Floer, Helmut Hofer, and Dietmar Salamon, *Transversality in elliptic Morse theory for the symplectic action*, Duke Math. J. **80** (1995), no. 1, 251–292. MR1360618 (96h:58024)
- [14] Kenji Fukaya, *Morse homotopy, A^∞ -category, and Floer homologies*, Proceedings of GARC Workshop on Geometry and Topology '93 (Seoul, 1993), Lecture Notes Ser., vol. 18, Seoul Nat. Univ., Seoul, 1993, pp. 1–102. MR **1270931** (95e:57053)
- [15] ———, *Morse homotopy and its quantization*, Geometric topology (Athens, GA, 1993), AMS/IP Stud. Adv. Math., vol. 2, Amer. Math. Soc., Providence, RI, 1997, pp. 409–440. MR **1470740** (98i:57061)
- [16] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono, *Lagrangian intersection Floer theory - anomaly and obstruction*.
- [17] Kenji Fukaya and Yong-Geun Oh, *Zero-loop open strings in the cotangent bundle and Morse homotopy*, Asian J. Math. **1** (1997), no. 1, 96–180. MR **1480992** (99e:58038)
- [18] William Fulton, *Introduction to toric varieties*, Annals of Mathematics Studies, vol. 131, Princeton University Press, Princeton, NJ, 1993. MR1234037 (94g:14028)
- [19] Alexander B. Givental, *Equivariant Gromov-Witten invariants*, Internat. Math. Res. Notices (1996), no. 13, 613–663. MR **1408320** (97e:14015)
- [20] Alexander Givental, *A mirror theorem for toric complete intersections*, Topological field theory, primitive forms and related topics (Kyoto, 1996), Progr. Math., vol. 160, Birkhäuser Boston, Boston, MA, 1998, pp. 141–175. MR **1653024** (2000a:14063)
- [21] M. Gromov, *Pseudoholomorphic curves in symplectic manifolds*, Invent. Math. **82** (1985), no. 2, 307–347. MR **809718** (87j:53053)
- [22] V. K. A. M. Gugenheim, *On the multiplicative structure of the de Rham theory*, J. Differential Geometry **11** (1976), no. 2, 309–314. MR 0418083 (54 #6127)
- [23] Allen Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002. MR **1867354** (2002k:55001)
- [24] Kentaro Hori and Cumrun Vafa, *Mirror Symmetry*, available at [arXiv:hep-th/0002222](#).
- [25] David G. C. Handron, *The Morse Complex for a Morse Function on a Manifold with Corners*, available at [arXiv:math.GT/0406486](#).
- [26] Maxim Kontsevich, *Homological algebra of mirror symmetry*, Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994), 1995, pp. 120–139, available at [arXiv:alg-geom/9411018](#). MR1403918 (97f:32040)
- [27] ———, *Lectures at ENS Paris, Spring 1998*, Notes by J. Bellaïche, J.-F. Dat, I. Marin, G. Racinet and H. Randriambololona.
- [28] Maxim Kontsevich and Yan Soibelman, *Homological mirror symmetry and torus fibrations*, Symplectic geometry and mirror symmetry (Seoul, 2000), World Sci. Publishing, River Edge, NJ, 2001, pp. 203–263. MR **1882331** (2003c:32025)
- [29] Naichung Conan Leung, *Mirror symmetry without corrections*, Comm. Anal. Geom. **13** (2005), no. 2, 287–331. MR **2154821** (2006c:32028)
- [30] J. Peter May, *Simplicial objects in algebraic topology*, Chicago Lectures in Mathematics, University of Chicago Press, Chicago, IL, 1992. Reprint of the 1967 original. MR **1206474** (93m:55025)
- [31] J. E. McClure, *On the chain-level intersection pairing for PL manifolds*, Geom. Topol. **10** (2006), 1391–1424 (electronic). MR 2255502
- [32] Dusa McDuff, *Symplectic manifolds with contact type boundaries*, Invent. Math. **103** (1991), no. 3, 651–671. MR **1091622** (92e:53042)
- [33] Grigory Mikhalkin, *Decomposition into pairs-of-pants for complex algebraic hypersurfaces*, Topology **43** (2004), no. 5, 1035–1065, available at [arXiv:math.GT/0205011](#). MR2079993
- [34] John Milnor, *Lectures on the h-cobordism theorem*, Notes by L. Siebenmann and J. Sondow, Princeton University Press, Princeton, N.J., 1965. MR 0190942 (32 #8352)
- [35] Robert Morelli, *The K-theory of a toric variety*, Adv. Math. **100** (1993), no. 2, 154–182. MR **1234308** (94j:14047)

- [36] Yong-Geun Oh, *Fredholm theory of holomorphic discs under the perturbation of boundary conditions*, Math. Z. **222** (1996), no. 3, 505–520. MR **1400206** (**97g**:58024)
- [37] Wei-Dong Ruan, *The Fukaya category of symplectic neighborhood of a non-Hausdorff manifold*, available at [arXiv:math.SG/0602119](https://arxiv.org/abs/math.SG/0602119).
- [38] Matthias Schwarz, *Morse homology*, Progress in Mathematics, vol. 111, Birkhäuser Verlag, Basel, 1993. MR **1239174** (**95a**:58022)
- [39] Paul Seidel, *Graded Lagrangian submanifolds*, Bull. Soc. Math. France **128** (2000), no. 1, 103–149, available at [arXiv:math.SG/9903049](https://arxiv.org/abs/math.SG/9903049) (English, with English and French summaries).MR1765826 (2001c:53114)
- [40] ———, *Fukaya categories and Picard-Lefschetz theory*, Zurich Lectures in Advanced Mathematics, European Mathematical Society (EMS), Zürich, 2008. MR 2441780
- [41] James Dillon Stasheff, *Homotopy associativity of H-spaces. I, II*, Trans. Amer. Math. Soc. **108** (1963), 275–292; *ibid.* **108** (1963), 293–312. MR 0158400 (28 #1623)
- [42] Andrew Strominger, Shing-Tung Yau, and Eric Zaslow, *Mirror symmetry is T-duality*, Nuclear Phys. B **479** (1996), no. 1-2, 243–259. MR **1429831** (**97j**:32022)
- [43] R. W. Thomason, *The classification of triangulated subcategories*, Compositio Math. **105** (1997), no. 1, 1–27. MR **1436741** (**98b**:18017)
- [44] Scott O. Wilson, *Geometric Structures on the Cochains of a Manifold*, available at [arXiv:math.GT/0505227](https://arxiv.org/abs/math.GT/0505227).

MASSACHUSETTS INSTITUTE OF TECHNOLOGY, 77 MASSACHUSETTS AVE., CAMBRIDGE, MA, USA

